Decentralized SSI Governance, the missing link in automating business decisions - v2

Authors: Rieks Joosten (TNO)
Scott Perry (SSPCPA)
Sterre den Breeijen (TNO)
Drummond Reed (Evernym)

Publication date v1: January 6th, 2021
ToIP draft approval: April 1st, 2021
Publication date v2: January 6th, 2022?

Keywords: "accreditation", "argument construction", "assurance community", "assurances", "certification", "credential catalog", "decision making", "decision support", "governance", "qualified data", "ssi infrastructure", "SSI-AC", "validation", "verification"

Acknowledgements: Thanks to Mary Lacity (University of Arkansas) and Savita Farooqui (SymSoft Solutions) for reviewing this document, and to the EU eSSIF-Lab project (EU H2020 Research and Innovation Programme - Grant Agreement Nº 871932) for partially funding this work.

Abstract

This paper explores some issues businesses might encounter as they decide to adopt SSI and reap its benefits, and need to transform their IT to accommodate for that. We assume an 'SSI infrastructure' exists, i.e.mechanisms and technology that enable such businesses to easily obtain 'qualified data', which is data that satisfies their (subjective) requirements for being valid for using it for their intended purposes. Our contribution is that we (1) identify several such issues, (2) suggest possible ways for addressing them and (3) propose to leverage already existing communities for governing parts of such solutions and providing assurances businesses need.

Summary

This whitepaper explores a set of related ideas that we collectively refer to as "decentralized SSI governance". The purpose of such governance is to support organizations as they transform their IT, their business-process artifacts, such as forms, and their policies to reap the benefits of SSI. This paper introduces SSI Assurance Communities (SSI-ACs) and identifies three specific governance topics: credential-types, accreditation and decision tree support. Tools and services are suggested that help with these topics. Furthermore, a distinction is made between what the business primarily cares about (business and business applications), and the technology and other things that are just expected to work (which we call "SSI-infrastructure").

Here are the main takeaways:

- Self-sovereign identity could save time and money on bureaucratic processes (form filling and validation - which includes login forms) by automating business decisions.
- Business decisions require governance to ensure the data used for making them qualifies (is valid) for that purpose, and continuously remains compliant with laws, regulations, and company policies (which regularly change).

- Companies could share this governance burden by supporting communities that support the provisioning of SSI-related assurances for their specific purposes.
- Such (focused) assurance communities could be supported by tools like credential catalogs and accreditation credentials.

We hope this paper inspires SSI proponents to develop not only real SSI infrastructure, but also assurance mechanisms for exchanging qualified data, i.e. data that comes with assurances that the party that processes such data requires in order for that processing to be ‘valid’ (well-grounded, justifiable, logically correct, ...) from that party’s perspective.

Table of Contents

Why SSI - Saving time and money on bureaucracy

SSI promises big benefits: better data quality, faster and cheaper data validation and decision-making, higher conversion rates and customer satisfaction, less churn, fewer IT-links[1], and operational costs, etcetera.

We estimate that in the Netherlands alone, monetary benefits are over 1Bn euro[2], waiting time for information can be reduced from hours (days, weeks) to seconds (minutes), and many IT-links can be dismantled.

SSI allows individuals to store and exchange qualified data when they want to, making it easier to fill in (digital) forms. This improves the individual's interaction experience with the system.

Currently, there are many problems that users can experience. For example, they do not understand what is being asked for, where to get the data or documents, and are challenged and/or experience errors when entering data into the form. Moreover, they experience the challenge to physically go somewhere to get a (signed) document, or scan documents and upload PDFs. The Dutch National Ombudsman has shown that everyone - including people with academic degrees, or lots of IT experience - faces these challenges.

On the other hand, organizations can also benefit from the verifiability components of SSI. At this moment, organizations need to deal with unqualified data (e.g. typed or uploaded, without any confidence about accuracy and provenance), errors and unsatisfied customers. SSI can provide them with the assurances[3] (e.g. regarding the provenance and integrity of such data, required accreditations of the data source, etc.) they need to qualify the data for the purposes they intend to use it.

Indirectly, yet nonetheless important for governments, is that SSI may help to reduce the 'digital divide'.[4] For example, people in lower socio-economic classes tend to find filling in forms more difficult, yet have to fill out most of the forms as they apply for social benefits and support. Typically, these forms are not the easiest to understand. As a result, these people give up, often for very good reasons, and therefore do not get benefits they are entitled to. Moreover, they do not file complaints if things go wrong. With SSI technology as their companion, all this can change for the better.

---

[1] IT-links connect IT-systems that are governed/managed by different parties (internal or external to an enterprise or government). It usually requires the setup and maintenance of a business contract, implementation of 'connectors' to convert data, authentication/authorization mechanisms, and a technical communications channel.

[2] This is a rough estimate that TNO did a few years ago. While there is no solid underpinning of these figures, representatives of various organizations (public notaries, banks, insurers) agree that this is a lower bound. Whether or not the figure turns out to be actually correct is less important than the fact that there is a wide consensus that we're talking big savings here.

[3] In this document we use the noun 'assurance' in the meaning of "anything (tangible) that a party considers to represent a contribution to the mitigation of a risk", which nicely aligns with the OED definition "A positive declaration intended to give confidence; a promise". In SSI contexts, this would typically be a (set of) claim(s), proofs etc., (to be) embedded in a credential or presentation.

[4] Representatives of various governmental bodies in the Netherlands have identified this as a concern they take quite seriously.

The challenge: automating business decisions based on qualified data

While most business decision makers can see the benefits and even want to reap them, acceptance is still an issue. Since SSI is still in its infancy, we would only expect 'innovators' and 'early adopters' to be interested.[5] So what are the issues being raised for not engaging now?

First, the technology still is not sufficiently mature. While many vendors are already developing SSI components and/or solutions, they do not yet (easily) interoperate with business applications or SSI components of other developers. Standardized specifications for APIs, credential exchange protocols, and other fundamental processes are still lacking.

Fortunately, these technical issues are being addressed and we expect to see them being resolved over the next two years. As a result, we foresee a generic SSI-infrastructure that is accessible for all (SSI-enabled) business applications, equally pervasive as the Internet IP-infrastructure. The main difference from a business point of view between the two is that the IP-infrastructure can be used to exchange *any* data, while the SSI-infrastructure will specifically be used for providing, requesting and obtaining *qualified* data.

Qualified data: data that comes with assurances[3], e.g. regarding its provenance and integrity (immutability), such that it qualifies as 'valid to be used for specific purposes of individual parties'.[6]

Figure 1: Qualified data exist in different types or formats.

[5] This would be in line with Rogers' Diffusion of Innovations theory.
[6] This has two perspectives. One is the usage perspective: whether or not data is 'qualified' depends on the assurances the user requires in order to be able to use it. This not only depends on the party that uses the data, but also the purposes for which it uses the data. The other is the provisioning perspective. Whether or not data 'qualifies' for use cannot be determined by the provider, but a provider can add assurances that increase the likelihood of it being used as such for increasingly more purposes by increasingly many users.

Second, business managers need to create and maintain the (machine readable) policies for the SSI infrastructure to provide such services. A party that requests and obtains data from such an infrastructure can provide a policy that specifies the kinds of credentials that are needed for its different business transactions, the sources (issuers) it should come from, etc. A party that provides (its) data through the infrastructure can specify the kinds of credentials it is willing to issue, and for each of them set conditions to determine whether or not to issue a credential and what data it should contain. Also, policies could give requirements to holders, on how to use specific credentials, etc. These policies can be hierarchical. For example, small policies of small businesses need to fit into the bigger policies of the jurisdiction the business is in.

This puts a burden on business governance processes: creating and maintaining the (machine readable) policies for issuing, storing, verifying and validating credentials wasn't required before. Reliable processes are required to develop and maintain such policies that are clear, unambiguous, complete, consistent, coherent and precise. This goes specifically for machine-readable policies, as machines cannot deal with incompleteness, inconsistencies, incoherence, impreciseness or ambiguities the way humans (often) can.

Finally, many network participants do not want to take the risk of being or becoming non-compliant. This means that the new SSI IT, the processes that use them and all related (digital) policies, have to be made explainable to third parties such as auditors (who are expected to have sufficient prerequisite knowledge) and decision makers.

Towards SSI Adoption and Transformation

Given the power it has to advance trustworthiness of claims, we expect that more organizations will adopt SSI and transform their business processes and IT systems when it becomes easier for them to do so.

For now, let's assume that the technology issues of interoperability, scalability etc. have been addressed, and that SSI infrastructure technology is a commodity, in the same way as the Internet IPv4/IPv6 infrastructure is currently. This would mean a set of standardized protocols exist (analogous to HTTP and FTP) for issuing credentials and obtaining them, as well as for requesting and obtaining 'presentations', i.e. a data constructs that contain (parts of) different credentials (issued by various parties) and (cryptographic) proofs of provenance, immutability, etc.[7]

Also, let's assume applications exist from various vendors that use these protocols, similar to e.g. web clients/browsers and servers that use HTTP. We already see 'user wallets' being offered that focus on obtaining credentials and constructing presentations upon request. Wallet vendors compete on UX, security, features, etc., but the basic functionality is the same. In a similar fashion, we expect to see ('issuer') components that provide parties with the capability of issuing credentials, and ('verifier') components that provide them with the capabilities to request and receive presentations, and to verify such presentations, i.e. evaluate whether a credential or presentation is an authentic and timely statement of the issuer or presenter respectively.

However, data that is verified is not necessarily also valid for the purposes that the receivers want to use it for. A credential that is issued by ACME Inc., that states Alice is over 18 may pass the verification checks, but it does not mean that everyone has to accept it as valid for their purposes.

[7] This is the mission of the Trust over IP stack from the ToIP Foundation. https://trustoverip.org/. Work is being done on coming to grips with (parts of) such protocols, such as DIF's Presentation Exchange protocol.

While verifiers obtain assurance that ACME issued a credential through SSI technology, it is also equally critical that verifiers have confidence in ACME's internal processes to issue that credential consistently and conformant to its own issuing scheme. Whenever a party receives credentials or presentations, it must not only verify them, but also validate them, i.e. determine whether or not the received data is fit for the purposes for which that party wants to process them.

There is an essential difference between verification and validation, which is that verification does not depend on the party that wants the data to be verified, whereas validation does. Before a party uses data to do some processing and/or decision making with, it should determine whether or not that data is valid for such purposes. Or at least, determine that the risk it runs associated with processing of possibly invalid data is acceptable.

From the business perspective, i.e. the 'data processing' perspective, an easy-to-use validation functionality is called for, one that evaluates data according to the requirements of the party that wants to process this data after it has been verified in a manner, and that leads to the determination of whether or not it is qualified for the purposes that party collected it. In short: a functionality that produces qualified data.

Adding this functionality to existing technology enables the construction of what we will call an 'SSI-gateway', which is a component that provides parties with an interface that enables them to:

- issue arbitrary data-sets, to which that gateway will then add proof of provenance, integrity and any other assurances prior to sending it out, according to a policy[8] that this party has specified for such kinds of data-sets. In effect, this means that a party can transmit a data set, and the gateway takes care of adding assurances, signatures etc, effectively turning the 'plain data' into 'qualified data'. Such qualified data may come in different forms, e.g. as Verifiable Credentials (VC), X.509 attribute certificates, SAML tokens, Attribute-Based Credentials (ABC), OpenID Connect scopes, etc.
- request a data-set of a specific kind/type that it needs for a specific purpose, and that comes with proof of provenance, integrity and any other assurances that the party needs for such purposes, according to a policy that this party has specified for such kinds of data-sets. In effect, this means that a party can request for qualified data, outsource the validation to the SSI-gateway, and obtain a (plain, but qualified) data set that it can immediately start to process.

Having such gateway functionality helps organizations to reduce the time and effort spent on validating information to the bare minimum. It allows them to focus on the contents (payloads) of credentials, the 'qualified data', rather than the 'envelopes' that the various kinds of credentials specify. This is similar to TCP/IP, the main interest of which is sending and receiving (unqualified) data, without being bothered by the 'envelope' (the IP-header).[9]

However, organizations *do* take an interest in, and pay attention to

- the actual (kinds of) qualified data (credential payloads), insofar they may serve purposes of an organization's business or its information processes;
- the kinds of qualified data that the organization itself might create and issue credentials with, insofar that fits with its business strategy/purposes;

[8] Such policies are typically provided by the issuer's governance framework.

[9] Extending this analogy, there might well be a business cases for 'SSI-providers', that, in analogy to Internet providers, provide their customers with SSI-gateways (internet routers/gateways) that it can use (after appropriate 'configuration' (policy specifications)) to provide (issue), store (hold), request and obtain (verify/validate) qualified data.

- assurances that it needs to designate data as being qualified, i.e. assurances that help mitigate any (unacceptable) risk that the organization perceives to run as a result of using such data in its information processes;
- how to construct arguments using this qualified data, insofar such arguments lead to qualified decisions (i.e. decisions that are based on qualified data) that are relevant to the organization - irrespective of whether such decisions are to be made by the organization itself or some other party.

Let's look at each of these more closely.

Three perspectives on qualified data: issuer, validator, and holder

As with credentials, qualified data can be looked at from different perspectives: the 'issuer perspective', the 'validator perspective'[10] and the 'holder perspective'. In the following sections, an explanation is given of what these perspectives are and how the different roles can act in a decentralized governance.

The Issuer Perspective

In its issuer role, a party is interested in creating value from sharing the knowledge it has about (other) entities (people, organizations, or things). For example, a government may decide to create credentials with citizen data (not just name and address, but also marital status and children, data concerning taxes, various permits, ownerships, guardianships, mandates, etc.). This may result in savings on bureaucracy that outweigh the costs of implementing credential issuing.

In order to reap such benefits, issuers should provide data that others will actually use, and do so with proofs of provenance, integrity and perhaps other assurances. This means that the issuer must communicate (advertise) the existence of such data in such a way that others can not only find it, but also decide whether or not that data is beneficial (for *them*) to use. An 'advertisement' not only needs to say 'what' the data is about, but also what its characteristics are (e.g. that this issued data is guaranteed to be 1-1 equal to the registrations of the issuer, unless the credential in which it is contained has expired or has been revoked), the liability (if any) that the issuer is prepared to take, conditions of use, etc.

From the issuer perspective, it isn't all that important what the 'envelope' is in which the data is conveyed, as long as assurances for provenance and integrity (and perhaps some others) for the data are in place. Verifiable credentials are trustworthy, but so are X.509 attribute certificates, Attribute-Based Credentials (ABCs), etc. Parties (i.e. organizations or individuals) that issue credentials may want the ability to specify which of these (not) to use, but that's a secondary concern.

In short, in its role as an issuer, a party is interested in creating value from sharing knowledge about other entities. A party's "issuing governance process" is concerned with making (and consistently reviewing and updating) decisions about e.g.:

- the kinds of qualified data it is willing to provide (what they consist of, what characteristics are to be ensured, liability to take, etc.;
- the kinds of credentials ('envelopes') it is willing to use for providing that data;
- under which conditions such credentials may be issued (e.g. only to a party that is mentioned in the qualified data);

[10] In the section on the validator perspective it is explained why we do not call this the 'verifier role'.

- the kinds of assurances that apply to issued credentials (e.g. by stating an assurance level, displaying certificates the issuer has that pertain to its issuing process, etc.).
- how all this is communicated: published, advertised and marketed (both in machine-readable and human readable form, for different purposes/audiences).

The Validator Perspective

Parties that seek data for processing (to obtain results that serve a specific purpose, e.g. to make a specific decision), need to validate such data. This is not simply ensuring that data has a correct syntax and semantics, but also ensuring that the data has other relevant properties that ensure the results of their processing are valid, i.e. fit for the purpose they are intended to serve. For example, weather data that is older than a week, while semantically correct and pertinent for an earlier time period, should not be used to compute tomorrow's weather forecast. Similarly, border agents would consider passport data that has expired invalid for the purpose of border crossing.

In 'traditional' SSI contexts, the role of 'verifier' is well-known, while a 'validator' role does not exist. We argue that there is a specific need for this role, as verification[11] and validation[12] are very different. Verification of a credential entails checking its signature and other proofs (determining its authenticity), whereas validating that same credential is determining (deciding) whether or not its data may be used for a specific purpose. Typically, the verification doesn't depend on context and hence can easily be automated. Validation on the other hand does depend on context, and specifically on the purposes for which the data is used. For example, consider a passport. Verifying a passport entails determining whether or not the document is authentic (not a fake), that there is a picture of its bearer, attributes such as name, date of birth, expiration date, a 'signature' of the issuing agency, etc. You can verify a passport without knowing the purpose(s) it may be used for. If you want to use your passport to enter a government building in the Netherlands, your passport will also be validated, which consists of making sure the passport hasn't expired for more than five years. However, when you apply for a visa at a Chinese embassy, your passport is only considered 'valid' if it expires more than 3 (or so) months after your projected entry date in China.

Validating data for use in specific contexts may require additional claims, which must be validated for the purpose of validating the original data, and this recursion may go on for some time.

In its role as a validator, a party validates data, i.e. it determines whether or not data qualifies as being valid for the purpose(s) it is to be used. This means that a party must inventory its arguments, (e.g. the business-rules that it will need for making particular decisions, algorithms for other kinds of processing etc). Such arguments will typically contain variables (i.e. placeholders) that need to be substituted with data/values at runtime after which the argument can be 'executed' (processed, evaluated). For every such argument, the validator must state the conditions (i.e. validation criteria) that the data must satisfy such that execution of the argument does not produce invalid results. Such conditions may distinguish between (in)valid data depending on its origin, its timeliness, its relevance for the purpose it is going to be processed, etc.

---

[11] The Verifiable Credentials Data Model document defines 'verification' as "*the evaluation of whether a verifiable credential or verifiable presentation is an authentic and timely statement of the issuer or presenter, respectively. This includes checking that: the credential (or presentation) conforms to the specification; the proof method is satisfied; and, if present, the status check succeeds*". It doesn't say anything about whether or not that statement can be used by a party in a way that is valid.

[12] The Verifiable Credentials Data Model document defines 'validation' as "*The assurance that a verifiable credential or a verifiable presentation meets the needs of a verifier and other dependent stakeholders.*" and subsequently declares it out of its scope, for the obvious reason that different stakeholders have different criteria for deciding what assurances they require, as they are in different situations and run different risks that these assurances serve to mitigate..

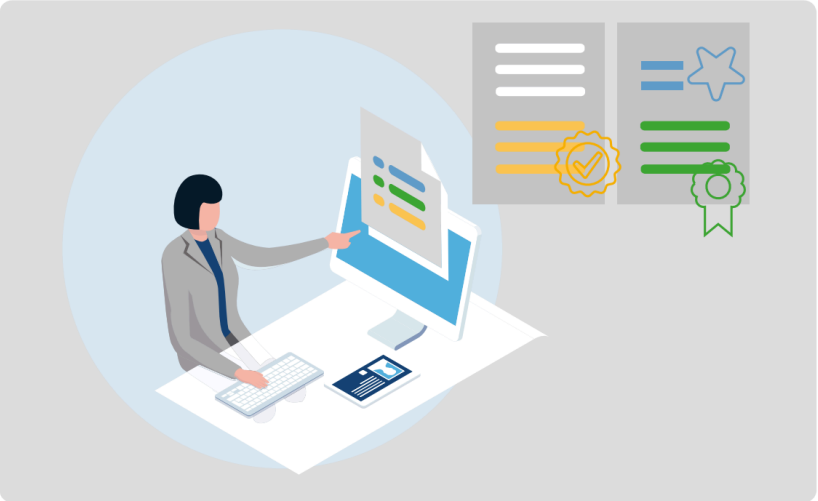

Figure 2: When designing a form, a party must state what information is requested and the conditions under which this data can (or cannot) be validated.

As a first step, a party can assess the communications/advertisements from parties that issue such data. This helps to determine the meaning (semantics) of that data, evaluate assurances provided, and other factors needed to determine its validity.

Having selected acceptable data sources enables the party to formulate requests[13] for data ((attributes in) credentials) which they need to evaluate their various arguments. Obtaining the data to evaluate a certain argument may require multiple requests, and data from multiple sources. Also, while requests would cover the data needed to evaluate the corresponding argument, it might also cover additional data that would provide additional assurances needed at runtime. For example, if a party would require data sources to comply with the requirements of some standard, it could formulate its requests to also ask for a certificate that proves such compliance to be valid at the time of the request.

Parties should establish new, or extend existing processes by which they not only govern the arguments for the various business decisions they need to make, but also the validation criteria that data must satisfy to ensure that when they are used to evaluate the argument, the risk that its conclusion is invalid is acceptable to them.

In short, as a validator, a party seeks to obtain data that is ‘valid’ when used for the purpose(s) it was obtained. In order to acquire/maintain the ability of validating data, it must establish and maintain a process that is concerned with making (and continually reviewing and updating) qualified decisions about e.g.:

- the mapping between variables in formulae on the one hand, and fields from credentials of specific issuers on the other;
- validation criteria for each of these mappings.

Doing this may cause the party to specify additional arguments that it needs to determine the validity of data used in other arguments, which also need to be governed.

[13] In technical SSI contexts, such requests are also known as ‘presentation requests’.

The Holder Perspective

In the holder role, a party is seeking to create value by collecting and managing (qualified) data for later use, by itself, or by presenting it to others e.g. within the context of a business transaction. The value comes mainly from the fact that handling such data is done electronically, using the SSI infrastructure, thereby avoiding most of the (costly, time consuming and annoying) problems that people face if they had to do this by hand (as explained earlier).

In the holder role, a party is seeking how to respond to presentation requests that it receives from other parties it transacts with. These other parties collect this data for the purpose of (populating formulae that, when evaluated, lead to them) making qualified decisions. Such requests would state which fields of specific kinds of (possibly multiple) credentials are needed, and might also state the validation criteria. Collecting the requested (meta-)data, wrapping it into a presentation format and 'issuing' it to the requester suggests that the party is actually performing in an issuer role. Note that collecting and presenting data in the requested format does not mean that the data needs to be changed.

A party in the holder role that responds to presentation requests (likes to) simultaneously fulfill the validator role, because the exchange of credential data is part of (a negotiation about) a business transaction. Users provide credentials as an equivalent of filling in forms. Few realize that this is only one half of the negotiation of a transaction, and that users, too, might want to learn more about this party that they provide their data to. If they can, they might be able to better determine whether it is a legitimate business or a scam, whether the party is properly accredited to handle the data (e.g. a health organization), etc. So, parties in a holder role should typically be enabled to asynchronously/simultaneously also perform in the validator role so as to balance the transaction negotiation.[14] Note that this implies that wallet apps should accommodate for this (by adding capabilities to send presentation requests and handling the responses), as should web servers (by adding capabilities that can handle presentation requests).

From this, it follows that parties that seek to engage in electronic transactions with one another using SSI technologies, must prepare to play the roles of issuer, holder and verifier/validator. in an asynchronous/simultaneous fashion. A party will take the issuer mindset when it issues credentials (or presentations, which are rather similar), the role of holder as it collects and stores received credentials (or presentation requests), and the role of verifier/validator as it issues presentation requests and receives and processes responses (presentations).

Automated decision making

As the Issuer and Validator perspectives show, the core of SSI is about the (electronic) exchange of data between parties, that each of them can use for its own purposes, and that often involves decision making of some kind. For example, if you want to rent a car online, you need to provide data that allows the rental car company to decide whether or not it provides you with a rental. And conversely, you may want to obtain data that allows you to decide whether or not the rental car company's offer is what you want, leading to a comparative analysis amongst rental car companies. For contracts to be executed, all parties involved in a transaction need compelling data to decide whether or not to commit.

This section deals with the requirements for supporting such decision making in an automated fashion.

[14] Some might recognize this as the 'verify the verifier' expression from CCI Use-Case 11.

Argument construction

We take ‘decision making’ to mean “the acceptance of one (qualified) proposition (the ‘conclusion’) on the basis of a set of other (qualified) propositions (‘premises’)”, according to the (business) logic used by the party that makes the decision. We will use the term ‘argument’ to refer to the set of premises and the business logic on which the conclusion is based; we use the term both for situations where premises can actually be evaluated, or (still) contain ‘variables’, i.e. placeholders that need to be replaced with actual data before it can be evaluated.[15]

Here are some examples of how we specify arguments in this document (the texts between `[` and `]` are the variables).

- “[authenticated person] owns [bank-account]” is an argument with variables that, when evaluated, results in a boolean value (‘true’ or ‘false’).
- “*rj2002050@gmail.com* owns *NL18RABO0123459876*” is the argument where variables have been substituted with actual values, and can be evaluated to produce ‘true’ or ‘false’.
- “([accountholder] owns [bank-account]) AND ([authenticated person] is mandated by [accountholder] to access [bank-account])” is an example of a more complex argument with variables.

For other purposes, e.g. machine processing, arguments may be specified in alternative syntax.

A well-designed argument is characterized by the fact that it is as simple as possible for making the decision it is used for. A ‘qualified argument’ is an argument that comes with (either the specifications for, or the actual) assurances that a party needs to reduce the risk of making an erroneous decision to an acceptable level.

> *Qualified argument*: Argument of a party that comes with (specifications for, or actual) assurances that reduce the risk of making an erroneous decision to an acceptable level.

In the examples above, the variable [authenticated person] was substituted with the value ‘*rj2002050@gmail.com*’. However, it is not at all certain that this value represents a person, let alone that this person was authenticated. In order for the argument to be considered qualified, it is not sufficient that the variables have been substituted with values, but also that somehow there is assurance that the risk that ‘*rj2002050@gmail.com*’ does not (properly) represent an authenticated person, is low enough to be acceptable.

Designing and establishing the policies that specify which arguments are to be used for what kinds of decisions is perhaps *the* core element of a governing party’s governance framework. Such argument specifications can take the form of a tree, where the root-node represents the final conclusion, and other nodes represent sub-arguments that their parent node combines using a function such as ‘AND’ or ‘OR’.

This breaking down of the argument into sub-arguments, in a tree-like structure, makes its application more flexible. For example, deciding whether or not a person should be given access to some digital service may be broken down in an argument for deciding this in case the person is an employee and the case in which (s)he is a customer.

[15] This is a common figure of speech: we can say “I am going to buy a bottle of wine” and 5 minutes later “I have bought a bottle of wine”. In the first case, ‘bottle of wine’ works as a variable, referring to a still unknown member of a class, and in the latter case it is a value that represents the actually existing element of that class.

Governance for arguments and decision making

The governance of arguments may be done in a distributed fashion, meaning that the task of establishing some (set of) sub-arguments can be assigned to the people that are best suited for that. For a bank, the simplest argument to decide whether or not someone may access a particular bank-account might be "[authenticated person] owns [bank-account]". There may however be legal situations in which someone else should also be given access. This can be done by introducing the argument "[authenticated person] has a legal right to access [bank-account]". We leave the governance of the argument that establishes whether or not this is the case up to the legal department.

The ability to break down arguments and distribute their governance can also help mitigate the risk of making wrong decisions. In high-risk situations, the governance of specific sub-arguments can be delegated to (e.g. risk and compliance) officers that have the skills to construct these arguments such that they provide the necessary assurance to reduce the risks to an acceptable level. An example is given in the section 'Tools and Services That Support Decision Making'

For completeness sake, we mention that a distinct part of the governance of arguments is the specification of the assurances that must be in place for data to qualify as valid for the position where it is applied in the argument. This has already been covered earlier, when we discussed validation policies. Basically, qualified data (for some argument) is data that comes with the assurances that are specified in the appropriate validation policy.

Assurances

A party that makes decisions in the context of (electronic) transactions would typically use arguments, parts of which are there to ascertain whether or not the risk-levels involved are acceptable. For example, a party that sells liquor might identify 'selling under age' as a risk it needs to mitigate. In an individual transaction, it could request that the customer prove he is older than the legal age limit. Such a proof, which could come as a claim in a credential, would then serve as an 'assurance' for mitigating this risk.

We will use the term 'assurance' to refer to anything (tangible) that a party considers to represent a contribution to the mitigation of a risk. We use the term particularly in the context of a party involved in negotiating and/or executing an electronic transaction. For example, a party that wants to buy a house could perceive the risk of it being like buying a pig in a poke to be mitigated if he received a clean inspection from an authorized house inspector. Another example would be a downpayment showing earnest commitment to the transaction.

What's important to realize here is that judgements pertaining to what constitutes a risk, the assessment of the associated risk-level in various situations, and what qualifies as an assurance, are all subjective, and will differ depending on the risk assessor's context. Before committing to a transaction proposal, a party will typically need multiple risks to be mitigated, in such a way that the net resulting risk is lowered to an acceptable level.[16] The decision of whether that is the case would typically require a related argument to be evaluated where variables or even sub-arguments can be considered such assurances. In the context of SSI, assurances can come as credentials, proofs, claims within credentials as well as in other forms. Here are some examples.

[16] It is not like every risk needs to be acceptable. For example, if a party considers the risk of a transaction going sour to be unacceptable, but the expected gain is high enough, he may nevertheless commit to it.

Our first example deals with ‘certification’ and ‘accreditation’. Prerequisite for both is a set of requirements, specified e.g. by a governance framework or standardization organization (e.g. ISO, NIST, ASC) that parties may want to be certified against. Certificates of compliance can be issued to this party by another party, usually called the ‘certifying party’. In SSI contexts, such certificates can be issued as credentials. Certifying parties (that issue such certificates) may also be accredited, i.e. certified against a set of requirements that ensure that parties that fulfill such requirements have proficient knowledge of the area in which they work and certify others, live up to high standards of integrity, etc.

Another example is generically referred to as ‘Level of Assurance’ or LoA. Prerequisite for this is the specification of:

- a generic question/concern, e.g. ‘How certain can I be that a user is who she claims to be?’
- one or more kinds of contexts (e.g. situations, equipment), that the question may be answered for, e.g. ‘users that attempt to login into some system’,
- a (limited) set of criteria that can be used to distinguish between situations/equipment in which the concern is addressed to a certain extent (level),
- a set of ‘levels’, (typically integers) one for each set of criteria.

The idea here is that it is much easier to convey a set of criteria constituting a ‘level’, and to have it processed by machines, than communicating the individual criteria that are being satisfied, even though some details/nuances are lost in the process.

LoAs are typically defined by standards organizations, such as NIST (NIST 800-63 which provides levels of assurance for identity proofing (IAL), authentication (AAL) and strength of an assertion in a federated environment (FAL)), or ISO (ISO/IEC 29115 provides a framework for managing entity authentication assurance, including 4 LoA’s for entity authentication).

Whenever a credential is issued that contains attributes pertaining to a context for which LoAs are defined, the issuer may choose to claim a certain LoA in (the payload of) that credential. And just as any other claim, other parties may or may not trust such claims to be true.

A very different kind of thing that might serve as an assurance is a cryptographic proof, e.g. as may be used in credentials or credential presentations. One can use such proofs e.g. to verify the immutability of a credential (i.e. that it hasn’t been changed since it was issued), or that all credentials in some set have been issued to the same party, or that the party to which a credential has been issued is 18+ years old.

A combination of assurances might constitute a new, stronger assurance. For example, the certainty of belief by a party that a public-key certificate is valid and that its corresponding private key can only be used by the owner registered in that certificate, combined with a digital signature that can be verified by that public key, constitutes a stronger claim over the provenance of that signature. Proofs may be used to establish links between credentials, and much more. Good cryptographic proofs have a solid mathematical basis that gives them the property that they cannot be denied.

While assurances can originate from different forms, SSI technology has equally many ways in which they can be requested and provided. Claims, proofs and even certificates can be embedded in the payloads of credentials and presentations. This allows parties to request not only the data they need for the transaction itself, but also to request assurances that they need. For example, a party that is electronically negotiating a transaction with another party, and that wants to mitigate the risks of that party using a rogue digital agent, may request an accreditation credential (issued

by a well-known/trusted accreditor) that states the set of security requirements that this agent satisfies.

SSI Assurance Communities (SSI-ACs)

The accompanying difficulties with the governance of information processes is a major obstacle for organizations to adopt SSI and transform their business processes and IT, particularly since trust is best achieved through a combination of machine-readable policies and regulations and human assessed variables. This is nothing new. Since computers have been used by businesses, the existence of this human-computer gap has been recognized, and people have tried to come to grips with it.

This document proposes several additions to the tools that have up till now been used for this. These tools assume the existence of an operational SSI infrastructure. Also, they capitalize on existing mechanisms that accept that trust and assurance work best in a community of parties that have some common objectives and work together to achieve them rather than having to do all the work individually.

Such parties do not necessarily seek to provide rules/standards that should be followed world-wide, but rather they consent to a set of rules within their stated scope of applicability. They will often allow others to join if they find that beneficial. We will use the term 'SSI Assurance Community' or SSI-AC to refer to such communities.[17]

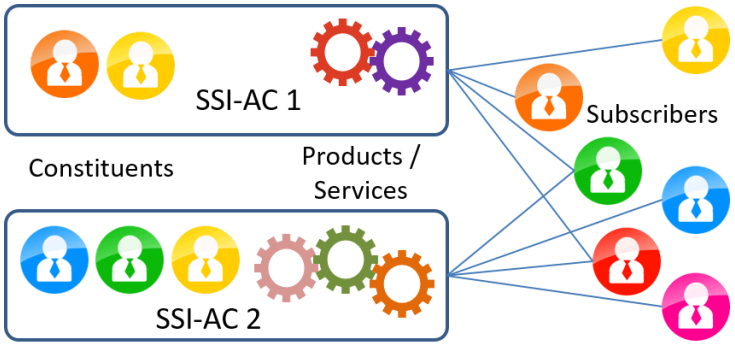

Figure 1: Constituents and Subscribers of an SSI-AC

An illustrative example of such SSI-AC could be the Dutch "Outbreak Management Team" (OMT) that consists of people from the National Institute for Public Health and the Environment, research institutions, medical societies, etc. The OMT is a standing organization that only convenes when "existing scenarios offer too little guidance" (to the government, the public, or enterprises regarding the ways in which to behave in situations with specified health risks) in the outbreak of an infectious disease. It is their task to propose additional guidance for specific situations. During the COVID-19 crisis, their advice to the Dutch government conveyed what citizens should (not) do, what kinds of enterprises are to be temporarily shut down, etc. This task made them ideally suited to also become a SSI-AC that specifies the necessary SSI support that organizations and citizens need, e.g. the kinds of credentials that citizens should be issued after having been tested, accreditation credentials for trusted testing facilities, etc.

[17] The term "SSI-AC" roughly corresponds to the combination of what the Sovrin Glossary and ToIP call a governance authority (SSI-AC Constituents) that serves on behalf of a trust community (SSI-AC members). What we propose in this paper is a more generalized form of this concept that goes into more details.

A SSI-AC is constructed and maintained by (representatives of) one[18] or more organizations - which we will call its 'constituents', and which form the governing party (sometimes also referred to as a governing authority as this body performs the governance). The SSI-AC (governing body) specifies its scope (what (not) to consider), and the products and services it will govern (and perhaps also realize) for the purpose of facilitating organizations to obtain assurances and build or maintain trust.

A SSI-AC may typically provide products and services that organizations need to adopt SSI and transform their business processes and IT, such as:

- Credential Type support, e.g. in terms of specifying their structure and semantics, supporting the advertisement of such specifications and the searching/finding of them;
- Accreditation support, e.g. in terms of maintaining lists of 'trusted issuers/verifiers/holders', specifying accreditation schemas and providing for the issuing and use of accreditation credentials associated with such schemes,
- Decision-Tree support, e.g. defining arguments for certain types of decisions, specifying validation policies for the data that such arguments need, providing a service that can push updates of such arguments and validation policies to subscribers, and perhaps even a generic (locally deployable) service for the evaluation of such arguments in given contexts.

Parties (i.e. organizations as well as individuals) that use products and/or services of a SSI-AC are called 'subscribers'. Parties are self-sovereign in deciding which SSI-AC(s) to subscribe to; they may also want to subscribe to multiple SSI-ACs[19], mirroring the meshed net of digital commerce. We use the term 'SSI-AC member' as a party that engages with the SSI-AC, either as a constituent or a subscriber.

We like to think of a SSI-AC as a service-oriented community that has a limited focus, such as in the example of the Dutch OMT, exists for the benefit of its members, and realizes that it is just one of many such communities. The latter means that a SSI-AC may decide to become a member of another SSI-AC, reaping the benefits associated with that SSI-AC. For example, the Dutch OMT may decide to become a constituent of some European Covid-19 SSI-AC, together with similar SSI-ACs from other EU countries. There, they could support credential types, accreditations and decision-tree support that would allow e.g. for EU cross-border use of credentials as they are governed by the different constituents.

In other words, a SSI-AC is a formal or informal, temporary or persistent organization that consists of different constituents (individuals, enterprises, governments) whose task is to at least govern[20] the SSI-AC (and optionally providing one or more of the products/services that it governs). This includes:

1. **define and maintain its scope**, i.e. the set of credential-types, jurisdiction(s), and domain(s) within which the SSI-AC aims to function, and the objectives it aims to pursue. This helps the Assurance Community to remain focused;

[18] GLEIF may be considered a single party that governs an SSI-AC.
[19] The discovery of SSI-ACs is a non-technical topic that resembles the way you discover the kinds of laws and regulations that apply to you(r business). You do that, e.g. by talking to peer organizations, governmental bodies (e.g. chamber of commerce), business associations, company lawyers, etc. It is conceivable that a SSI-AC exists, or will be created, that will maintain a register of SSI-ACs (that may be accredited as such by some scheme).
[20] See also the COVID-19 Credentials ("C19C") Governance Framework.

2. **define the services and products that the SSI-AC governs**[21], examples of which are given above;
3. **define restrictions, artifacts etc. that are necessary for the provisioning of such products and services.** Restrictions include e.g. liabilities of issuers under the SSI-AC, specific choices in semantics, credential mappings (i.e. stating the list of alternative credentials that can be used if a certain credential is being requested), defining credential lifecycle, defining audit processes, etc.
   Artifacts include e.g. the mechanism/process scheme by which it is determined if an organization qualifies as a 'trusted issuer'[22] within the SSI-AC;
4. **the operational details of producing such products and services**, e.g. which organizations will, or are allowed to perform these operations, endpoints of e.g. CI/CVs where products and services can be obtained, etc.

Figure 3: A party can use a decision tree as a basis for decision making, not only when taking the issuing, validating or holder role, but also as (part or member of) a SSI-AC.

[21] This does not necessarily mean that the SSI-AC itself provides such products/services; it may also outsource this.
[22] The idea is that verifiers may decide to trust any issuer in the trusted-issuers list of particular SSI-ACs, effectively (partially, at least) outsourcing the vetting work of what would be trusted issuers.

Tools for Supporting SSI-ACs

In this chapter, we present a vision for tooling that allows SSI-ACs to support SSI in the short, medium and long term.

In this chapter we make the following assumptions:

- The sovereignty of parties implies that each of them can issue any kind of credential. Other parties, also when they are referred to as an authority of some kind, cannot deprive parties of this ability. Therefore, we do not elaborate on what are called 'trusted issuers'. Rather, we provide support for accreditation against arbitrary accreditation schemes, which can be used to address the 'trusted issuers' issues, as well as several other/related issues.
- A credential consists of an 'envelope' and a 'payload'. The payload is a set of statements that parties issue (and other parties are interested in). The 'envelope' is just another (yet possibly standardized) means of transporting that payload with some basic assurances. For example, a bare passport document (as it comes from the printing house) is an 'envelope', and the payload is whatever is added (statements, photo, fingerprints) etc. The VC data model spec specifies an envelope where the payload/contents are to be put in the `credentialSubject`[23] section. There are many more kinds of envelopes, e.g. X.509 attribute certificates, attribute-based certificates, OpenID tokens, etc.
- Parties are typically interested in the payload (and they assume that the envelope is properly 'managed' by the infrastructure). They may choose to (not) use certain kinds of envelopes for issuing a payload, or to (not) accept certain kinds of envelopes. The policies for specifying such preferences to the (commodity) SSI infrastructure that we assume will exist are out of scope for this paper.

Short Term: Tools for Credential Markets

In the short term, we will need tools that support credential markets as well as related governance. Such tools will enable parties to search for and find the kinds of credentials that other parties have on offer, such that they can decide they are suitable to be used for (a) specific purpose(s). Conversely, such tools also enable parties to advertise the kinds of credentials they are willing and able to issue, to specify the assurances that they come with as well as any other data that their 'customers' may need. We will elaborate on two tools we think will be indispensable, but they will undoubtedly be complemented with others.

Credential Catalogs

We use the term '**Credential Catalog**' to refer to a functional component, or an operational service, that parties can use to advertise the kinds of (payload contents of) credentials that they are willing and capable of issuing. The purpose of such 'advertisements' is to help a party find the kinds of credentials that are actually available, and also help it evaluate their usefulness for the specific objectives that such a party may have - specifically, for using them in an argument for a specific decision. This means that such advertisements should allow for ways to advertise not only syntax and semantics, but also about other characteristics that parties may need e.g. to determine whether or not (and if so, under which conditions) credentials can be used in arguments for making specific kinds of decisions.

[23] This may be a bit confusing because 'subject' usually refers to some entity, whereas 'credentialSubject' is a set of statements about different entities.

The exact nature of the information that parties may want to publish in payloads will depend on the (needs of their) prospective users. It will contain a (JSON, XML or other) schema that specifies which statements are in the payload, and what they mean (semantics). Also, it would specify the kind(s) of 'envelope'(s) that may be used.

Other information might include the process that the issuing party has followed to verify the data that it puts in such payloads, standards or regulations that it has followed, constraints for use, pricing/payment mechanisms, applicable accreditations, liabilities it is willing to accept, etc.

A very simple credential-catalog (proof-of-principle) was built for the Odyssey hackathon (November, 2020), for the purpose of experimenting, and the elicitation of further functional requirements.

SSI-ACs may use credential catalogs to document and advertise the types of payloads that they govern. They can use it as a platform

- to inform parties that may want to join the SSI-AC, or subscribe to its services,
- to store, or link to documentation about appropriate accreditation schemes,
- to inform verifiers how it can determine which parties have been accredited for some accreditation scheme (see section 'Supporting Accreditation Schemes']),
- to inform parties how they can get certified against such schemes, the kinds of payloads they can then obtain that certify this, etc.

Yellow Pages service

We use the term '**Yellow Pages service**' to refer to a functional component, or an operational service, that parties can use to search for and discover payload-types that other parties may govern and/or issue (in a variety of 'envelopes').

Typically, such a service would allow various kinds of search mechanisms, e.g. allow searches based on specific contents (e.g. addresses, or names), assurance levels, issuers, analogues, keywords/categories etc. A very simple yellow pages service (proof-of-principle) was built for the Odyssey hackathon (november, 2020), for the purpose of experimenting, and the elicitation of further functional requirements.

Parties that operate a credential catalog and those that run a yellow pages service will find ways to cooperate, as it is obvious that in general, both benefit from each other's existence. However, the specifics of the credential catalog services and yellow-pages services, such as the expected needs of their respective customers (e.g. the kinds of credentials, or kinds of characteristics they will be looking for), will determine what arrangements between them will be beneficial, and impose requirements e.g. on the credential type advertisements. This is an area for further market-research.

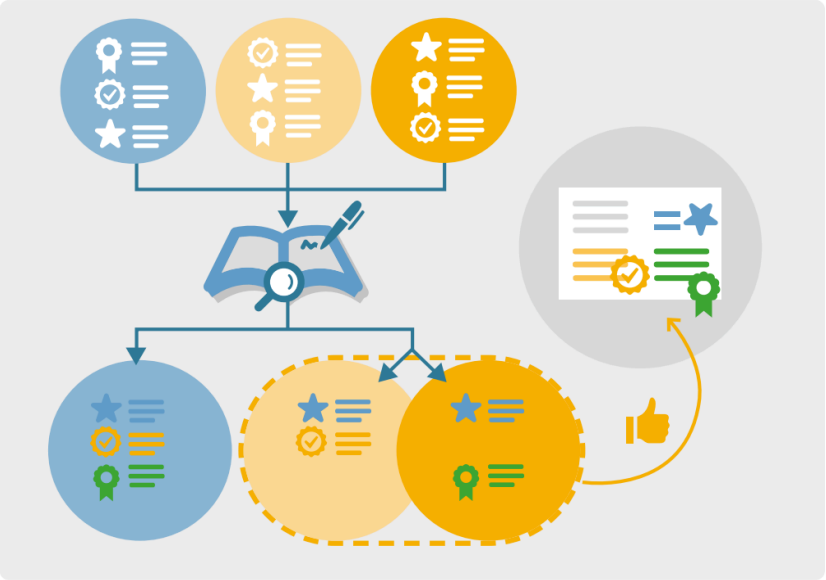

Figure 4: Parties can describe the different credential types that they are willing to issue in a credential catalog. The other side of this functional component can then be used by (other) parties to look for pay-load types of credentials, as a yellow pages service.

Medium Term: Supporting Accreditation and Certification

In the medium term, we will need support for (decentralized) accreditation and certification.

An obvious application is the issuance of credentials that contain a certificate stating that its subject is successfully accredited or certified against some existing scheme. A party that is certified, e.g. against ISO 9001 or ISO 27001, could embed the credential that states this into any credential that it issues, which might serve as an assurance for relying parties.

We also expect to see SSI-ACs developing their own derivatives of existing accreditation and certification schemes, as well as new schemes that have been created for purposes that are specific to the SSI-AC.

Let's look at an example for this. In the section on the validator perspective we explained that a party makes decisions by evaluating associated arguments that have been populated with validated data[24], i.e. data that it considers valid to be used in such arguments. It is up to the individual parties to decide what data is (in)valid for what kinds of decisions. Its "validator" governance process is concerned with making (and continually reviewing and updating) such decisions.

Data is valid (to be used in some kind of argument) if the party 'trusts' it. That is to say: if the party has come to believe it is true. Trust is not a simple 'yes' or 'no' - it is an analog scale, a 'trust level', that increases as data comes with assurances. But in the end, a 'yes'/'no' decision must be made about whether or not to accept data as valid, i.e. whether or not the trust level has passed some 'clutch point' beyond which the perceived risk of using invalid data, that remains in the face of the assurances, is acceptable.

[24] 'validation' is the process that parties use to determine whether or not a specific data element can be used in a specific argument. These processes are not only subjective - they are unique for every party, but they also depend on the kind of argument that is under consideration. Data that a party considers valid for one kind of decision may not be valid for another, and vice versa.

SSI-ACs may play a significant role in providing a specific kind of assurance, namely ‘accreditations’ and ‘certifications’. There’s a subtle difference between the two.[25]

Typically certification is a documented attestation (by the certifying party) of the conformity of a product, process or service to a set of requirements that are specified in some (certification) scheme.

Accreditation is the formal recognition of a party’s competence to conduct a specific activity such as issuing, holding and/or verifying a specific kind of credential, or certifying parties. Accreditation typically comes with a documented attestation (by the accrediting party), that states that the party has demonstrated compliance with a set of requirements that are specified in some (accreditation) scheme.

Thus, a successful certification or accreditation results in a written statement of conformance, and is signed by the certifying/accrediting party. Such signed statements are easily checked and may serve as assurances for parties that trust the certification/accreditation process.

SSI-ACs are in an excellent position to define both certification/accreditation schemes and the acceptable certificate/accreditation artifacts. Specifically, they could define them for purposes that are specific to the community, and serve the (shared) purposes of its members (and non-members that see the benefits of using this).

There’s nothing new here - everything already exists, and SSI-ACs can readily use the existing practices in the existing communities. What may differ is the kind of artifacts that SSI-ACs can associate with being accredited. Since SSI-ACs operate in the ‘SSI world’, it makes sense to link accreditations with (verifiable, or other kinds of) credentials, or a registration in a ledger or database, which is also easily checked.

Accreditation Credentials

Let’s consider the use of credential payloads for accreditation and validation purposes. We use the term ‘accreditation credential’ to refer to a credential (payload):

- for which a SSI-AC has established and published a credential type specification (using its Credential Catalog);
- that states (implicitly and/or explicitly[26]) at least:
    - the party to which the accreditation credential has been issued;
    - the accreditation scheme whose requirements have been fulfilled by that party[27];
    - the credential-types to which the accreditation credential applies;
    - the (cryptographic) proof methods and associated data that allow the proofs to be verified for the assurances that the accreditation scheme specifies. This obviously

[25] See What is the difference between accreditation and certification? (pbctoday.co.uk).

[26] Implicitly: by describing in the accreditation credential type specification what this trust consists of, i.e. what it has decided to believe about the (issuer) organization, and what assurances it has obtained (and perhaps also: how it has obtained these assurances) that convinced it to hold this trust/belief.
Explicitly: by including claims in the accreditation credential that state the obtained assurances for this particular issuer organization.

[27] The accreditation scheme implies the functions that the party can be trusted to properly execute. So you might have a ‘trusted issuer’, or ‘trusted verifier’ accreditation schemes for specific kinds of credentials, as well as certifying parties against such schemes.

includes the proof of provenance and integrity (signature), as well as other proofs that may be required under the accreditation scheme[28];
- the party that has audited the accredited party[29], the date of the audit, and perhaps some other audit-related attributes;
- … (etc.)

Parties can be accredited for different functions, against different schemes[30]. We may have such schemes not only for accrediting 'trusted issuers', but also for 'trusted verifiers'. A party that has been certified against a trusted verifier scheme of a SSI-AC would be trusted by other parties (to the degree of their confidence in the SSI-AC), to request and use credentials of specific kinds only for specific purposes.[31] Similarly, IT services and equipment (i.e. wallet equipment[32]) may be certified, instilling confidence in parties that they will interact with 'trusted verifiers' and/or 'trusted issuers' according to the requirements of that certification scheme.

If SSI-ACs were to provide accreditation credentials, this would be of enormous help for their subscribers to limit the number of credentials that otherwise would need to be requested. A party that needs credentials for making a particular decision, and that can simply request that the issuing party must have a specific (issuer) accreditation credential of some SSI-AC, doesn't even have to know the name of the issuing party. For example, a party that requests a credential from a user that states the result of a Covid-19 test does not need to know about each and every (accredited) test-lab; it only needs to know that the credential was issued by a party that has the corresponding accreditation by the appropriate SSI-AC.

Trustworthy Credentials of a SSI-AC

We introduce the term 'Trustworthy Credential (of a SSI-AC)' for a credential,

- whose type is specified by (and published in the Credential Catalog of) that SSI-AC, thereby implying that it deems the credential issuing scheme trustworthy;
- of which the 'envelope' (metadata) includes one or more accreditation credential payloads, one of which is issued by, or on behalf of that SSI-AC, and is of a type that is specified (and published) by that SSI-AC.

The fact that such credentials contain an accreditation credential of the SSI-AC means that it comes with the particular assurances as stated in the associated accreditation scheme. The

---

[28] Doing this, and including the payload of the accreditation credential in every credential that is issued under this regime, enables verifiers to check the trustworthiness of the issuer based on the SSI-ACs assessment without needing to know who the actual issuer is. The exact/preferred ways of doing this remain to be determined. Candidates include using (pseudonymous) DIDs, and also with ZKP VCs.

[29] If a SSI-AC decides to outsource its accreditation process, it should make sure that the associated accreditation credential types that it specifies make verification and validation as easy as possible. There are several possibilities: the SSI-AC can allow organizations to use its 'accreditation credential signing service' if they present a credential that states they are a SSI-AC accreditor (so every accreditation credential is signed with a single key that is owned/controlled by the SSI-AC). Alternatively, accreditors may issue accreditation credentials that also include the accreditation credential of the accreditor (there is recursion here…). And there are more ways.

[30] As with ISO certification: different management systems of an organization may be certified: the quality management system (ISO 9001), environment management system (ISO 14001), the information security management system (ISO 27001), etc.

[31] In the Netherlands, the Burger Service Number (BSN - i.e. the Dutch social security number) may only be legally used within the government, for health purposes, and by banks. Certifying other organizations against a 'trusted verifier' scheme might provide sufficient assurances to allow them to use government issued credentials that contain this number.

[32] Note that where 'trusted issuers' and 'trusted verifiers' refer to parties (individuals, organizations), for the holder role we need the actual equipment (wallet app, edge/cloud agent) to be certified, because it is that equipment that will do the actual receiving of credentials, and creating presentations.

credential is ‘trustworthy’ for any party that appreciates the assurances provided by that accreditation scheme.

Example of a Trustworthy and Accreditation Credential.

Here is an example of what a simple Accreditation Credential, in the form of a VC, might look like. This credential (of type `DHACAccreditationCredential`) has been issued by a SSI-AC named `DHAC`. It has been issued to a party called ACME and asserts that ACME is a trusted issuer for credentials of type `DHAC:Covid19TestResult`.

```
{
  "@context": [
    "https://www.w3.org/2018/credentials/v1",
    "https://www.w3.org/2018/credentials/examples/v1"
  ],
  "id": "http://DHAC.org/accreditationCredentials/1872",
  "type": ["VerifiableCredential", "DHACAccreditationCredential"],
  "issuer": "https://DHAC.org/issuers/4",
  "issuanceDate": "2020-04-30T11:17:24Z",
  "credentialSubject": {
    "id": "did:example:ebfeb1f712ebc6f1c276e12ec21",
    "name": "ACME, Inc.",
    "trustedIssuerCredentialType": "DHAC:Covid19TestResult",
    "trustedIssuerProofType": {
      "type": "RsaSignature2018",
      "verificationMethod": "https://acme.com/issuers/keys/1"
    }
  },
  "proof": {
    "type": "RsaSignature2018",
    "created": "2019-06-18T21:19:10Z",
    "proofPurpose": "assertionMethod",
    "verificationMethod": "https://DHAC.org/issuers/keys/3",
    "jws": "eyJhbG...GHSrQyHUdBBPM"
  }
}
```

Figure 2: Simple example of Accreditation Credential

Figure 3 shows an example of a Trustworthy Credential as it might have been issued by ACME to a person called Wayne Dodge. It states that Wayne has been tested on April 30th, 2020, and that the result of the test was negative[33]. It also contains ACME’s accreditation credential that allows a verifier to obtain assurance that, according to SSI Assurance Community DHAC, ACME is a trusted issuer for this credential.

[33] This credential content is fictitious, and may be replaced with whatever else may be appropriate.

```
{
    "@context": [
      "https://www.w3.org/2018/credentials/v1",
      "https://www.w3.org/2018/credentials/examples/v1"
    ],
    "id": "http://acme.com/credentials/DHAC_Covid19TestResult/172",
    "type": ["VerifiableCredential", "DHAC:Covid19TestResult"],
    "issuer": "did:example:ebfeb1f712ebc6f1c276e12ec21",
    "issuanceDate": "2020-05-01T12:13:14Z",
    "credentialSubject": {
      "id": "did:example:2bdcc0b259683e194e48037ea21e15d3",
      "name": "Wayne Dodge",
      "covid19TestResult": {
          "tested": "2020-04-30T11:19:10Z",
          "result": "negative"
      },
      "accreditationCredential": [ {
        "@context": [
            "https://www.w3.org/2018/credentials/v1",
            "https://www.w3.org/2018/credentials/examples/v1"
        ],
        "id": "http://DHAC.org/accreditationCredentials/1872",
        "type": ["VerifiableCredential", "DHACAccreditationCredential"],
        "issuer": "https://DHAC.org/issuers/4",
        "issuanceDate": "2020-04-30T11:17:24Z",
        "credentialSubject": {
            "id": "did:example:ebfeb1f712ebc6f1c276e12ec21",
            "name": "ACME, Inc.",
            "trustedIssuerCredentialType": "DHAC:Covid19TestResult",
            "trustedIssuerProofType": {
                "type": "RsaSignature2018",
                "verificationMethod": "https://acme.com/issuers/keys/1"
            }
        } ],
        "proof": {
            "type": "RsaSignature2018",
            "created": "2019-06-18T21:19:10Z",
            "proofPurpose": "assertionMethod",
            "verificationMethod": "https://DHAC.org/issuers/keys/3",
            "jws": "eyJhbG...GHSrQyHUdBBPM"
        }
    },
    "proof": {
      "type": "RsaSignature2018",
      "created": "2019-06-18T21:19:10Z",
      "proofPurpose": "assertionMethod",
      "verificationMethod": "https://DHAC.org/issuers/keys/3",
      "jws": "eyJhbG...GHSrQyHUdBBPM"
    }
  }
```


Figure 3: Example of a Trustworthy Credential

Certification Credentials

In considering the use of credential payloads for certification purposes, we use the term 'certification credential' to refer to a credential (payload):

- for which a SSI-AC has established and published a credential type specification (using its Credential Catalog, and thereby implying that it deems the certification credential issuing scheme trustworthy);
- that states (implicitly and/or explicitly[34]) at least:
    - the specific entity, or entity-type[35], to which the credential content applies. Such an entity could be a wallet app, a cloud wallet, a proxy, etc.;
    - the certification scheme whose requirements have been fulfilled by that entity[36];
    - the kinds of entities to which the certification credential applies (e.g. wallets, cloud-agents, proxies, etc.;
    - the (cryptographic) proof methods and associated data that allow the proofs to be verified for the assurances that the certification scheme specifies. This obviously includes the proof of provenance and integrity (signature), as well as other proofs that may be required under the certification scheme[37];
    - the party that has audited the certified entity(-type), the date of the audit, the auditor's accreditation certificate, and perhaps some other audit-related attributes;
    - … (etc.)

Entities can be certified for different functions, against different schemes. We may have schemes not only for certifying wallet apps, but also for cloud wallets. Something that has been certified against a trusted wallet app scheme of a SSI-AC would be trusted, at least within the scope of that SSI-AC, e.g. to issue credentials to and/or exchange presentations with.

If SSI-ACs were to provide certification credentials, this would be of enormous help for their subscribers to reduce the complexity of specifying what credentials to request. Specifically, it would enable parties to deploy SSI technology that could verify that the technology of other parties that it is transacting with functions in accordance with the requirements of some certification scheme. In 'business terms': it allows a party to do automated checking of the integrity of the agents of the parties it is doing electronic transactions with. For example, for high-risk/high-value transactions, the party could decide to only commit if the equipment (say: cloud agent or wallet agent) of its peer-party comes with a specific certificate, that is issued by an accredited auditor and the certificate states the requirements the party needs its peers to have.

In the further future, this may be complemented with 'automated attestation services', i.e. IT-services to which equipment may be registered as soon as it is being deployed, and that can

[34] Implicitly: by describing this in the certification credential type specification. Explicitly: by including claims in the certification credential that state the certified characteristics of the entity.
[35] Often, it is not necessary to certify each individual entity. It may suffice if the entity is of a specific kind for which there is a certification scheme, and the party that has manufactured that entity is accredited to (only) produce instances of that type that comply with the requirements of hte certification scheme.
[36] The certification scheme specifies the requirements of the characteristics that the entity must have in order to be certified. So you might have a 'trusted wallet app', or 'trusted cloud agent' certification scheme, associated credential-types, as well as (accredited) parties that can (type-)certify such entities.
[37] Doing this, and including the payload of the certification credential in every credential that is issued under this regime, enables verifiers to check the trustworthiness of the issuer based on the SSI-ACs assessment without needing to know who the actual issuer is. The exact/preferred ways of doing this remain to be determined. Candidates include using (pseudonymous) DIDs, and also with ZKP VCs.

later, at any time, issue (ephemeral) certificate credentials that state that the state (integrity) of the equipment hasn't changed since it was registered. Such certificate credentials, together with manufacturer certificates, may provide great assurance.

Long Term: Tools and Services For The Future

In the long run, we expect to see more real-time functions that do not yet exist, but could have a significant impact on the uptake and (ease of) use of SSI. They will also have governance implications, and may end up in the list of topics that a governance process seeks to address. This section provides two examples of this.

Decision Support

A first example is intended to help members of SSI-ACs outsource those parts of their operational decision making that are not their core business. In this case, the core business of a bank is concerned with financial topics, and in their day to day operations; decisions need to be made whether or not a person can access some bank account and the maximum amount of money that can be transferred out of such an account, etc.

In this case, the challenge for governance is that there are situations in which the decisive part of the argument is not financial, but of a regulatory, legal (or other) nature. For example, a user that is the financial guardian of some person would have a legal right to inspect the bank account of its dependent, and transfer funds out of it.

In order to ensure that operational decisions take legal and regulatory considerations into account, the bank would need a process that continuously watches for changes in the laws and regulations of every jurisdiction that has some legal right. If a candidate change is detected, its personnel must identify the bank systems that are affected by this change, make impact assessments, and propose management decisions for dealing with it. Then, the actual change requests must be established, budgets and project teams allocated, the projects executed, tested and in the end: deployed. Accommodating for legal and regulatory changes takes significant amounts of time and resources. This problem will only grow over time, as the number of (changes in) applicable laws and regulations only increases.

However, arguments that lead to a decision are typically a construction of terms (sub-arguments or statements) that are 'glued' together with AND and OR statements, which allows them to be 'modularized' into (sets of) terms each of which address a particular concern. For example, we can write "[access] = [authenticated-user] AND [user-owns-bank-account]" which means that the access decision (to the bank account) requires [authenticated-user] (the argument that decides whether or not the user is authenticated) to be 'true' and also that [user-owns-bank-account] (the argument that computes whether or not the user is the/an owner of the bank account) is 'true'. Note that each such term may (recursively) require simple or complex arguments to be evaluated.

This property allows arguments to be easily extended with other terms. For example, if the bank decides to allow owners of bank accounts to permit others for accessing their bank accounts, it could specify an argument for the term [permitted-user] that expresses the condition that a permission exists for the bank account, where the owner of that account has created the permission, the user is permitted to exercise some rights, and the rights consist of 'accessing the bank account' and 'transferring money out of the bank account'. Once the argument of this term is

defined, the [access] decision is easily rewritten as “[access] = [authenticated-user] AND ([user-owns-bank-account] OR [permitted-user])”.

Extending an argument for an operational decision consists of two parts. First, the term by which the argument is extended must be defined. This requires knowledge about the topic or concern that the term needs to address. Secondly, at runtime, the argument must be evaluated to produce its result. This entails gathering the required (and qualified) data, and using it to compute the outcome.

An SSI-AC may support its subscribers in both parts. First, as an SSI-AC is typically organized around a set of objectives and topics that are of interest to every of its subscribers, it is realistic to assume there is expert knowledge about such topics in the SSI-AC that is also kept up-to-date. Thus, members of the SSI-AC may use such knowledge to create and maintain arguments for making various decisions for which such knowledge is required. For example, an SSI-AC around regulatory compliance of banks could create arguments for deciding whether or not a user has a right to access and/or transfer funds from specific bank accounts.

Second, an SSI-AC might provide policies (rules, working-instructions and other guidance) that its members could use for obtaining qualified data to evaluate such arguments. And in particular, as the SSI-infrastructure that we envisage becomes more and more available, such policies might be created in a machine-readable form, so that the necessary qualified data can automatically be requested and obtained.

Parties that subscribe to both the argument definitions and such policies could do away with the tedious tasks of keeping tabs on the changes in knowledge and doing whatever is necessary to make their IT-systems comply with the regulations. This saves them a lot of money and effort.

Obviously, the SSI-AC would then also need to have an ‘argument catalog’ in which it can advertise the kinds of decisions and associated arguments, the kinds of (qualified) data that are involved, the ways in which subscribers can access them, etc.

Cryptographically Enforceable (Issuer) Policies

We also expect to see a new function emerging that we will provisionally refer to as Cryptographically Enforceable (Issuer) Policies. It will enable parties to encrypt arbitrary plaintexts ‘under a policy’, i.e. a criterion that other parties may or may not satisfy. The resulting cryptogram can only be successfully decrypted by a party that actually satisfies this policy. It is based on a technique called Ciphertext-Policy Attribute Based Encryption (CP-ABE). Its application within SSI is currently being explored in a draft article. Here is a summary of what it can do.

Parties can use this to encrypt one or more parts of credentials they issue, which can then only be read by holders or verifiers that satisfy the associated policy. For example, if a person participates in a medical trial, she may be given a credential that states the details of her participation, but whether she gets the trial medicine or a placebo could be hidden from her, and anyone else that is not a member of the project team so as not to disrupt the trial.

Similarly, parties that create presentations can encrypt parts thereof such that the verifier can only use it if it satisfies the holder's policy. A consumer organization or other SSI-AC could define a policy that can distinguish between (il)legitimate webshops. Its subscribers can then use this policy to encrypt presentations with, and a webshop can then only see the data if it is legitimate.

This would then require SSI-ACs to provide support for KeySmith offerings, i.e. advertisements by KeySmiths that allow encryptors (issuers, holders) to select encryption keys that allow for encrypting under the policies they have specified, and decryptors (holders, verifiers) to obtain an (associated) decryption key. Also, parties may want to extend their advertisements, e.g., of credentials they issue, by documenting which parts of those will be encrypted and where decryptors may go to obtain an associated decryption key.

Conclusions

We have explored a set of related ideas that we collectively refer to as “decentralized SSI governance”, the main purpose of these concepts is to help organizations transform their IT, business-process artifacts and policies to enable them to use SSI and reap its benefits.

We have identified several of such benefits, not just saving time and money, but also that using SSI may contribute to diminishing the ‘digital divide’, because it prevents people from giving up on filling in digital forms as they encounter all sorts of difficulties.

We have also identified obstacles to the adoption of SSI by organizations, postulating that the electronic exchange of qualified data - using an SSI infrastructure - should be just as easy as the electronic exchange of arbitrary data - using the Internet, i.e. TCP/IP infrastructure. We have assumed that over time, such an SSI infrastructure will become available as a commodity and use the current trust networks as a basis.

The adoption and transformation challenges that remain have to do with

- **qualified data**, i.e. data that comes with assurances regarding provenance, integrity, and possibly others. We have described the related concerns from the perspective of providing such data (issuer perspective), obtaining and using such data for further processing (validator perspective), and holding such data (holder perspective). We conclude that organizations need governance processes that establish and maintain policies for providing and validating such data.
- **argument construction**, i.e. governance processes for establishing arguments, which is the reasoning by which decisions are being made in an organization, and the possibility to outsource the (partial) design of such arguments, particularly in cases where specific (expert) knowledge and experience is required.
- **assurances** related to qualified data, other than those regarding its provenance and integrity. We have identified various sources of such assurances as well as different ways to communicate them in terms of (un)qualified data.

Then, we have introduced the idea of an SSI Assurance Community (SSI-AC), that is a community of organizations that usually already exists, in which its members already work together for some purposes and have established trust mechanisms that they rely on. We propose to leverage such communities and their trust mechanisms for the purpose of furthering the adoption of SSI, and the related transformation of organizations.

We do so by proposing the development of tools

- **for the short term - credential markets**. One example is ‘credential catalogs’, that organizations and/or SSI-ACs can use to advertise the definitions of the payloads of credentials that they govern, providing all information that other parties may need to determine whether or not using such credentials would be beneficial for them. Another

example is the ‘yellow pages service’, which is a service that allows such parties to find the various credential catalogs that may be of interest to them.

- **for the medium term - accreditation and certification**. Existing schemes for certification and accreditation can be adapted and new ones that are specific to SSI can be established. Certificates of compliance can be issued as credentials, and we proposed a generic kind of accreditation credential of which we showed how it can be used in a way that is integratable with SSI infrastructure.
- **for the long term - tools and services**. We foresee new tools and services to be developed in the future, for which we gave two examples. First, we showed how SSI-ACs can support the operational decision making of their members by creating arguments, and policies for obtaining the qualified data that allows them to be evaluated. This is particularly beneficial for those parts of the operational decision making that require knowledge outside of the core business of such members, e.g. legal and regulatory.

  A second example is where parties are enabled to specify cryptographically enforceable policies that enable them to encrypt (parts of) the credentials or presentations they issue ‘under a policy’, i.e. in such a way that they can only be decrypted by parties that satisfy such a policy.

Taking it all together, SSI-ACs propose to help the adoption of SSI, so that the current bureaucratic information exchange processes are easier, both for individuals as well as organizations. Using the SSI-ACs, it is possible to leverage the present trust communities and corresponding mechanisms, while at the same time automating the business decisions that are experienced as cumbersome. Different credential types, accreditation, machine readable policies and decision tree support are an essential part in decentralized governance. To reach these potentials that the SSI-ACs can offer, different tooling and services are suggested. Such tools and services will need to be better specified, and some perhaps also standardized.
