

# SoK: Trusting Self-Sovereign Identity

Evan Krul

University of New South Wales  
Cyber Security Cooperative Research Centre  
Sydney, NSW, Australia  
e.krul@unsw.edu.au

Sushmita Ruj

University of New South Wales  
Sydney, NSW, Australia  
sushmita.ruj@unsw.edu.au

Hye-young Paik

University of New South Wales  
Sydney, NSW, Australia  
h.paik@unsw.edu.au

Salil S. Kanhere

University of New South Wales  
Sydney, NSW, Australia  
salil.kanhere@unsw.edu.au

## ABSTRACT

Digital identity is evolving from centralized systems to a decentralized approach known as Self-Sovereign Identity (SSI). SSI empowers individuals to control their digital identities, eliminating reliance on third-party data custodians and reducing the risk of data breaches. However, the concept of trust in SSI remains complex and fragmented. This paper systematically analyzes trust in SSI in light of its components and threats posed by various actors in the system. As a result, we derive three distinct trust models that capture the threats and mitigations identified across SSI literature and implementations. Our work provides a foundational framework for future SSI research and development, including a comprehensive catalogue of SSI components and design requirements for trust, shortcomings in existing SSI systems and areas for further exploration.

## KEYWORDS

digital identity, self-sovereign identity, identity data privacy

## 1 INTRODUCTION

Digital identity encompasses not only usernames and passwords but also our online presentation, interactions with others, and how the world perceives us [5]. In 2023, Louisiana passed a law requiring all websites containing adult content to verify the user's age. Louisiana's law mandated users present their state-issued digital identities, such as driver's licenses, to access these websites [39]. This is not a singular event; digital identities are becoming ubiquitous tools in many applications, including age verification, know-your-customer compliance [121, 127], and academic credential verification [146]. Digital identity has expanded beyond mere name and email addresses in an online account.

In the current ecosystem, digital identity is managed centrally through *siloes* (one identity per service) or *federated* (single sign-on) identity systems [129]. These centralized identity management systems typically rely on a single entity, such as a government agency or a private entity, to issue and manage identities. This concentration of power has created a single point of failure, making these

systems susceptible to data breaches and misuse. In recent years, there have been numerous high-profile breaches that highlight the inherent risks of centralized identity management [21, 23, 58, 118].

Furthermore, the reliance on centralized authorities has placed undue trust on these data custodians. Users are forced to relinquish control over their personal identity information, with limited transparency and accountability regarding how their data is used. This lack of control raises concerns about potential privacy violations, misuse of identity data, and the possibility of the central authority gaining access to sensitive information about the usage of an identity [73, 131].

In response to these concerns, a new identity paradigm has emerged: Self-Sovereign Identity (SSI), aiming to revolutionize identity management by empowering individuals to take control of their digital identity in the absence of a centralized authority [5]. The essence of SSI lies in the issuance and verification of user credentials. These credentials, issued by trusted entities, attest to the authenticity of a particular attribute of a user's identity. By holding the credentials, users eliminate reliance on third-party data custodians, reduce the potential for data breaches, and ensure they retain control over their identity.

While SSI bears the same underlying principle of fostering trust between unknown parties as previous systems [29], the extent of what "trust" means in this context has undergone considerable discussion. In both conceptual and implementation level discussions of SSI, there are disagreements on what trust is already available between participants and what trust needs to be built up through the design of the system [10, 72]. The discussions and summaries in this paper will illustrate how complex and fragmented the field of SSI has become. Without a clear agreement on the threats the participants face, there can be no guidance on what SSI must address in its design and requirements.

The paper aims to systematically analyze the elements of SSI in light of trust requirements and assumptions. We do this by breaking down SSI into its individual components to understand the underlying trust assumptions (or alternatively threats) that motivated their design requirements. This analysis provides insights into the design of an SSI system and identifies areas for further research and development. A detailed description of the data collection and analysis process is included in Appendix A.

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license visit <https://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.



*Proceedings on Privacy Enhancing Technologies* 2024(3), 297–313

© 2024 Copyright held by the owner/author(s).

<https://doi.org/10.56553/popets-2024-0079>

### 1.1 Contributions

We claim the following three major contributions:

- (1) A comprehensive analysis of trust in SSI, introducing three distinct trust models to capture the varied levels of assumed trust across literature and implementations. These models are built by identifying the components of SSI and the threats they face from various actors in the system.
- (2) An analysis of known implementations of SSI, both in academic literature and industry products, detailing their trust assumptions made and the implementation approaches taken to address the assumed threats identified in (1).
- (3) A foundational framework for future SSI research and development by providing a comprehensive catalogue of SSI components and design requirements according to trust models. We identify shortcomings in existing SSI systems and highlight areas for further exploration.

The remainder of this paper is structured as follows: Section 1.2 provides an overview of related work. Section 1.3 introduces the core concepts of SSI, laying the foundation for subsequent discussions. Section 2 uses these core concepts to establish a minimal trust model for SSI, highlighting the design requirements necessary to mitigate associated threats. Section 3 delves into the threats that issuers, identity owners, and service providers face, extending beyond the minimal trust model. Section 4 consolidates these threats into two additional threat models. Section 5 evaluates the state-of-the-art academic literature and industry implementations, and Section 6 concludes with a discussion of open research directions.

### 1.2 Related Work

We situate our work within the landscape of well-established surveys on SSI and closely related topics. Through six categorizations, we differentiate our contribution from existing literature and present a summary in Table 1.

**Academic Literature and Industry Implementations;** We present a comprehensive analysis of both academic literature and industry implementations. There have been several surveys that collect and analyze academic literature on SSI. Soltani et al. [131] performed a thorough survey and provided a detailed and in-depth survey of SSI. They provided compelling motivation for SSI and then introduced core components with a strong focus on academic literature, briefly considering a subset of industry implementations. Schardong and Custódio [126] extensively covered publishing venues and provided a taxonomy of academic literature. Zhu and Badr [154] described a sub-set of academic and industry contributions focused on blockchain identity for the internet-of-things. Mühle et al. [108] did not consider any industry contributions, only surveying academic contributions. Ernstberger et al. [51] and Bistarelli et al. [17] both broadly collected DID methods from academia and industry. Lim et al. [95] and Satybaldy et al. [123] evaluated a sub-set of industry blockchain identity systems. Grüner et al. [65] only evaluated a single industry implementation.

**Catalogue of Components;** This work presents a catalogue of components of SSI and their design requirements. Schardong and Custódio [126] independently collected a large catalogue, while Ernstberger et al. [51] presented a limited catalogue.

**Threats and Trust Assumptions;** A key focus of our contribution is trust in SSI and the associated threats. Hardman [70] presented the most complete SSI threat model focused on preventing credential fraud. Grüner et al. [65] provided a threat model for a single implementation, while Ernstberger et al. [51] had limited discussion on threats faced. Lim et al. [95] introduced a trust model for blockchain implementations.

**Foundation Framework;** Satybaldy et al. [123] is the only contribution presenting a framework to evaluate future work. They evaluated five SSI constructions with a framework inspired by [5].

**Not Limited to Blockchain Identity;** A key distinction with our contribution is that we do not present a survey focused on blockchain-based SSI contributions. Schardong and Custódio [126] were the only other work not to have a scope limited to blockchain identities.

Table 1: Related Works Comparison

	Our Work	Grüner et al. [65]	Soltani et al. [131]	Satybaldy et al. [123]	Zhu and Badr [154]	Mühle et al. [108]	Schardong and Custódio [126]	Lim et al. [95]	Ernstberger et al. [51]	Hardman [70]	Bistarelli et al. [17]
Academic Literature	●	●	○	○	○	○	○	○	○	○	○
Industry Implementations	●	●	○	○	○	○	○	○	○	○	○
Catalogue of Components	●	○	○	○	○	○	○	○	○	○	○
Threats and Trust Assumptions	●	○	○	○	○	○	○	○	○	○	○
Foundational Framework	●	○	○	○	○	○	○	○	○	○	○
Not Limited to Blockchain Identity	●	○	○	○	○	○	○	○	○	○	○

● included ● partially included ○ omitted

Although not a survey, the closest work to ours is Hardman [70] who introduced a credential fraud threat model in a blockchain-based Aries RFC. Our review of existing SSI surveys identified a lack of consensus on SSI components and their underlying design motivations. No existing work provides a comprehensive catalogue of SSI components, an in-depth analysis of trust assumptions, and an extensive collection of academic literature and industry implementations presented in this paper. Consequently, no existing work serves as a sufficient foundational framework for future research.

### 1.3 Overview of SSI and Terminology




Figure 1: SSI Credential Exchange

Self-sovereign identity (SSI) is founded on the issuance, storage, and presentation of identity credentials [56, 57]. Three distinct roles are involved in SSI:

**Identity Owners** require identification and collect credentials that represent their identities.

**Issuers** are authorized entities that issue credentials to identity owners (government registries issuing driver’s licenses).

**Service Providers** request credentials from identity owners to verify their identities.

To facilitate further discussion, we clarify the definition of an identity. Aligned with the Sovrin Glossary [57], this work defines identity as a collection of *attributes*, where each attribute represents a particular characteristic, such as name or age. Attributes are verified through attestations from issuers. These attestations are termed *claims*. A collection of claims compiled by an issuer to describe an identity owner is a *credential*.

Credentials are stored in a personal wallet, typically on the identity owner’s device. The exchange of credentials between parties is illustrated in Figure 1. Issuers share verifiable credentials with identity owners, who can then share verifiable presentations of these credentials with service providers. The additional term “verifiable” indicates that the credentials may be verified by the service providers; this concept is further discussed in Section 2.1.3. Service providers rely on the credential issuer’s authority without direct interaction. They infer the issuer’s identity from the credential and use their existing knowledge of the issuer and their authority to determine whether to accept the credential (only a government authority can issue a driver’s licence). This separation of credentials and isolation of the issuer from the service provider during credential presentations mitigate the concentration of issuer power that challenges the centralized identity models [5, 56].

This process of issuing, storing, and presenting identity credentials, known as the *credential exchange*, is often portrayed as the entirety of SSI [65, 126]. While it represents a crucial component of SSI, it does not capture all of the components. For that, we will turn to the *Trust Over IP Stack*.

## 2 TRUST OVER IP AND THE TRUSTFUL MODEL FOR SSI

The Trust Over IP (ToIP) Stack introduced by the *ToIP Foundation*<sup>1</sup> is a four-layer architectural framework designed to establish trust between users over the Internet and other digital networks. It is inspired by the TCP/IP stack, which standardized packet exchange and enabled the development of the Internet. The ToIP Stack, however, emphasizes human trust through establishing digital identities and secure communication.

We outline the structure of SSI through the trust layers of the ToIP stack to analyze fundamental concepts and components of SSI in the context of forming trust. Following the analysis, we introduce a minimal trust model, our baseline *trustful model*.

### 2.1 SSI through Trust Over IP Layers

Our adaptation of ToIP layers tailored to SSI with the appropriate terminologies [57] is shown in Figure 2. We walk through each layer, introduce the relevant components of SSI, and summarise implementation choices available for the components.

**2.1.1 Layer 1: Trust Support.** Layer 1 provides a way to identify participants and verify the provenance of these identities through trust anchors. This layer provides a trusted foundation upon which the subsequent layers are constructed.

**Identification.** Every party represents themselves with a globally unique self-generated identity handle [51]. Recipients communicating with the party use this identifier to derive public attributes

<sup>1</sup><https://trustoverip.org>




Figure 2: SSI through Trust Over IP Stack

(e.g., public key) [133]. Issuers create long-lived identifiers shared broadly, while other entities may create new identifiers for each communication. Two general implementation approaches to providing identifiers are decentralized identifiers (DIDs) or public keys.

DIDs are a URI of the format *did : method : identifier*. The DID *method* specifies how the identifier is de-referenced. De-referencing a DID produces a DID Document (DIDDoc), which details the entity’s public attributes. Every DIDDoc must include at least a public key for the subject [133]. Alternatively, a simple public key can serve as an identifier, eliminating the DID/DIDDoc abstraction<sup>2</sup>.

**Trust Anchors.** Trust anchors affirm the provenance of identities and public attributes (DIDDocs). They are “trusted components” that serve as starting points for establishing trust across the other layers [75]. For implementation, in distributed DID-based scenarios, verifiable data registries (VDRs, e.g., blockchain) often serve as trust anchors. However, public key infrastructure (PKI) may also be employed, notably when raw public keys are used instead of DIDs.

Recording a DIDDoc on a publicly accessible VDR, such as a blockchain, guarantees the record’s immutability and auditability [53, 97]. The method associated with a VDR-backed DID specifies the location of the corresponding DIDDoc [17]. For instance, *did : ethr* indicates the identifier is an Ethereum address that locates the DIDDoc [53]. Trust anchors are not always VDRs; *did : web* (and DID-less approaches) rely on traditional public key infrastructure to anchor identities and public attributes [62].

**2.1.2 Layer 2: Trust Spanning.** Layer 2 focuses solely on facilitating communication between agents and wallets of Layer 3 participants. Agents, the software libraries managing credential requests, facilitate this communication and interact with wallets [57]. For

<sup>2</sup>DID methods used for short-lived DIDs, where a trust anchor is not required, frequently require that the DID is the public key (*did : key* [101], *did : peer*, [45]). In these instances, DIDDoc resolution involves extracting the public key directly from the DID itself.

implementation, building upon DIDs (Layer 1), DIDComm [40] establishes a suite of protocols for inter-entity messaging. DIDComm utilizes the public keys associated with DIDs to encrypt messages exchanged between entities.

**2.1.3 Layer 3: Trust Tasks.** Layer 3 facilitates the credential exchange. It defines the standardized, interoperable credential formats (e.g., driver’s licenses). It also specifies verification mechanisms (e.g., signatures), credential management protocols (collection, storage, presentation), and optional extensions to the core process.

**Credential Formats.** Credentials adhere to a standardized format comprising credential metadata, claims, and proof. Credential metadata encompasses the issuer’s identity (e.g., DID), credential type (e.g., driver’s license), and issuance date, among other details [132]. Claims represent the identity attributes verified by the issuer (e.g., name, birthdate, address, licence no.). The proof ensures the credentials are tamper-proof for integrity and provenance. Presentations are temporary, ephemeral documents containing verifiable credentials that service providers request. Some works consider presentations synonymous with verifiable credentials [56]. Regarding implementation, W3C verifiable credentials are the predominant data model, employing JSON-LD to define the credential structure [63, 151]. The AnonCreds JSON format links credential definitions to DIDs stored on VDR (Layer 1) [79].

**Verification Mechanisms (Proofs).** Credentials contain a verification mechanism as a *proof* of the credential’s authenticity. This proof may be provided through a signature, where the issuer signs the entire credential (metadata and claims) with their private key. Service providers verify the credential’s authenticity by retrieving the issuer’s public key from their DID. For implementation, ECDSA signatures are a common choice [87]. All surveyed credential profiles use signatures as a verification mechanism, although some do not present the signature to the service provider; instead, the identity owner presents a zero-knowledge proof of knowledge of a valid signature [28, 141] (discussed further in Sections 3.2.1 and 3.2.2).

**Protocols.** Protocols define the procedures controlling the exchange of credentials, encompassing how identity owners request credentials from issuers, how service providers request presentations from identity owners, and how identity owners present credentials to service providers. Many proposals implement various protocols. For instance, multiple related projects from OpenID address the issuing and presentation of verifiable credentials [98, 140, 149].

**Extensions.** Extensions add additional SSI features to the credential exchange. Our discussion utilizes these extensions to introduce threats and analyze their risk and mitigations to derive various trust models (Section 3).

**2.1.4 Layer 4: Trust Applications.** Layer 4 depicts the identity applications ecosystem (e.g., government, health, banking). Each application is considered to have *Trust Assumptions*. They define trusted and untrusted parties, interactions, and components. For instance, an application may distrust identity owners and assume they may present fraudulent credentials.

*Deriving Trust Requirements from Trust Assumptions.* Intuitively, a trust assumption implies a threat – e.g., suggesting the credential presented to the service provider could be fraudulent.

In fact, we can derive that a trust assumption (or lack thereof) implies the application accepting the presence of a *threat*. Threats pose a *risk* to some participants and must be addressed through *constraints*. That is, a Layer 4 application that distrusts identity owners and assumes they may present fraudulent credentials mandates the implementation of tamper-proof credentials in Layer 3.

Based on this observation, in our analysis of trust requirements for trust models, we consider that the human trust assumptions made in Layer 4 are stated as threats, which defines *constraints* that need to be addressed by the components of Layers 1-3.

As a summary, Table 2 outlines common implementations or standardization initiatives for each SSI component discussed in this section. Trust assumptions in Layer 4 determine the use of these components.

**Table 2: Common SSI Trust Over IP Layer Implementations**

Layer 4	<b>Application</b>	Various trust assumptions impacting choices Layer 1-3
Layer 3	<b>Format</b>	W3C VC [132], JSON-LD [63], JWT-VC [103], SD-JWT-VC [139], AnonCred JSON [79], MDOC [147], ICOA DTC [84], x.509, IRMA XML [55]
	<b>Verification Mechanism</b>	ECDSA [87], RS256, CL [28], BBS [141], BoundBBS [145], zk-SNARK [114], GGM-Merkle [109]
	<b>Protocol</b>	CHAPI [31], WACI [8], Aries Protocols [77], OpenID4VCI [98], OpenID4VP [140], SOIPv2 [149]
	<b>Extensions</b>	<i>Detailed in Section 3</i>
Layer 2	<b>Communication</b>	DIDComm [40]
Layer 1	<b>Identification</b>	DID (indy [44], sov [97], ion [52], ebsi [48], key [101]), Any Public Key (x.509 [74])
	<b>Trust Anchor</b>	Blockchain (Indy DLT [83], Bitcoin [52], EBSI Trust Registries [49], Ethereum [53], Parity [138], Iroha [82], Fabric [80]) IOTA [54], GNU Name System [125], Any PKI (CA, FIDO [6])

## 2.2 The Trustful Model

Here, we introduce the first trust model: *trustful* SSI. Trustful SSI encapsulates the fundamental principles of SSI, establishing the *minimum responsibilities* for each component. As mentioned, the model is represented as a collection of threats accepted as posing risks under the application’s trust assumptions. As a minimal model, it considers the threats relating to potential credential integrity and external actors, along with mitigation strategies. To clarify the design and responsibilities of the SSI components, we group the threats by components as a functional unit that implements a group of features.

### 2.2.1 Identity Owner Identification.

**Threat 2.1.** If identity owners can present credentials without a verification mechanism, they may present credentials (e.g., a driver’s licence) that were never issued and do not represent their identity (e.g., the identity owner is not licensed to drive).

**Risk:** A service provider may authorize an identity owner who should not have been authorized.

**Constraint:** Credentials (and the presentations that encompass them) must be verifiable. Credentials must be issued alongside some cryptographic proof that the service provider can use to verify the integrity of credentials.

**Mitigations:** *Identity owner identification* is implemented through the verifiable credentials in Layer 3. When the proof portion of verifiable credentials uses a signature to ensure the integrity of credentials the minimum requirement from the signature is three basic algorithms that provide *key generation, signing, and verification* functions. These signatures must satisfy correctness and security [87] such that a service provider will always accept a valid signature by an issuer, and importantly, an adversary cannot forge a signature that passes the verification function but was not created by the issuer.

### 2.2.2 Issuer Identification.

**Threat 2.2.** An issuer may not have the authority to issue credentials of any type (e.g., may only issue state driver’s licenses).

**Risk:** A service provider may accept a credential issued by an entity that did not have the authority to attest to the identity claims.

**Constraint:** The issuer of a credential must be identifiable from a verifiable presentation such that service providers can be assured that the claims they have been presented were issued by an issuer that they have some *human trust* in (trust that a government entity may issue driver’s licenses).

**Mitigations:** Layers 1 and 3 work together to address this constraint. Issuers embed and sign their DID in the Layer 3 verifiable credentials. Service providers then use this DID to derive what entity issued the credential and verify their authority. The immutability of credentials (Threat 2.1) ensures that once issued, the issuer’s DID can not be modified. The DID is then de-referenced to the publicly recorded DIDDoc to verify the identity and rights of the issuer.

### 2.2.3 Non-Transferability.

**Threat 2.3.** The identity owner may present a correctly issued credential; however, it was not issued to them. Rather, it was collected through collusion with another identity owner or stolen from an unknowing identity owner (e.g., Alice presents Bob’s driver’s licence as if it was hers).

**Risk:** A service provider may authorize an identity owner by incorrectly believing claims that do not actually describe them.

**Constraint:** Verifiable credentials and their presentations must be tied to the identity owner such that they can not be transferred incorrectly to another identity owner.

**Mitigations:** The identity owner’s DID is included in the verifiable credential. The immutability (Section 2.2.1) ensures that the DID cannot be modified after issuance. Ownership of the credential is provided by proving knowledge of the secret key associated with the public key of the DID (e.g., a Sigma protocol [41]). The issuer ensures that the identity owner knows the secret key when they embed the DID in the credential. Service providers perform the same verification when receiving presentations. For an adversary to present a verifiable credential as their own, they must have knowledge of the secret key associated with the embedded DID.

### 2.2.4 Protected Communication.

**Threat 2.4.** An outside entity may intercept a credential in transit, either passively or actively. Passive interception would allow

the outside entity to learn about the contents of the presentation, while active interception would allow the outside entity to modify the contents of the presentation.

**Risk:** For identity owners, their credentials in transit may unknowingly have their private identity information stolen or modified.

**Constraint:** The credential exchange of Layer 2 must be secured to protect communication between entities from outside entities.

**Mitigations:** The encryption in DIDComm protocols (Layer 2) ensures that no eavesdropper may learn the contents of verifiable presentations. Additionally, credentials cannot be modified by a third party if they are encrypted in transit. This includes the threats of replay [136], and man-in-the-middle attacks [60] with protections such as token-binding in the Layer 3 verifiable presentations [100].

**Threat 2.5.** An outside entity may interfere with DID resolution, resulting in an incorrect DIDDoc being de-referenced.

**Risk:** A service provider may attempt to verify a credential using an incorrect DIDDoc of the issuer, resulting in either the authorization of an inauthentic or rejection of an authentic credential.

**Constraint:** Layer 1 must provide assurances in the provenance of resolved DIDDocs.

**Mitigations:** When de-referencing a DID, service providers must be sure to use protected communication channels. Additionally, the registries publicly hosting the DIDDocs should provide provenance guarantees that the DIDDoc has not been modified. At the very least, these registries need to provide some level of assurance specified by the Layer 4 application (e.g., blockchain-backed DID [53, 97], certificate authority-backed DID [62]).

## 3 EXTENDING TRUST REQUIREMENTS

While adequate in many scenarios, trustful SSI may not fully address the intricacies and security demands of all digital identity applications. This is evident in government-issued digital identities, where the identity data may be highly sensitive. The trustful model’s assumption of service provider trustworthiness overlooks potential misuse or insecure storage of identity data.

This motivates more cautious approaches to SSI, where a broader set of trust requirements should be considered. For instance, since applications that follow the *trustful* model do not consider the service providers a threat, identity owners may share their entire driver’s licence credentials to prove their age. This is problematic as a dishonest service provider could collect the additional claims in a driver’s licence for malicious purposes or collect them innocently but store them insecurely.

Here, we broaden the trust requirements by introducing more threats perceived by issuers, identity holders and service providers.

### 3.1 For Issuers

The *trustful model* did not include threats that impacted the issuers. Here, we introduce an extended set of threats for issuers.

**3.1.1 Identity On-Boarding.** New identity owners may be brought into an SSI system through a process known as on-boarding.

**Threat 3.1.** Potential identity owners may attempt to mislead an issuer when being on-boarded such that they are issued credentials that do not represent their true identity (e.g., being issued a driver's licence credential with an incorrect name or driver's class).

**Risk:** Issuers could be misled into issuing invalid credentials.

**Constraint:** Credentials should only be issued if some assurances exist regarding the validity of the identity being attested.

**Mitigations:** Much of the issuer's responsibility in identifying a potential identity owner is offloaded outside technical constraints. Issuers are expected to perform some level of due diligence [143]. Take, for example, a government-issued identity; it is reasonable to expect the government to have prior knowledge of the individual requesting a credential. Large portions of Sovrin's contributions in Layer 4 have been creating frameworks and specifications that define the responsible behaviour of identity issuers to protect the system's integrity by having only responsible issuers prepare credentials [56]. This approach can be augmented with frequent audits, such as what is done for certificate authorities in traditional public key infrastructure (PKI) [9]. Two optional extensions in Layer 3 provide solutions to address this threat.

*Verifying Existing Credentials;* Taking the service provider role, issuers may request a credential that verifies common claims present in both an existing credential and the new one being pursued. If the claims align with the claims expected by the issuer, they issue the requested credential (e.g., a name claim matches another claim).

To note, there is no indication within verifiable credentials that the issuer performed this verification. This approach must assume that the additional credential requested was not compromised through similar means. Finally, a potential identity owner may not have an existing SSI verifiable credential they can present, especially if they are a new user.

*eID Derivation;* Deriving identity claims from outside SSI by bootstrapping SSI to existing identity schemes allows for low-friction identity on-boarding. Abraham et al. [1] presented an SSI implementation that enabled deriving credentials from existing eID identity infrastructure. Identity owners created non-interactive zero-knowledge proofs [114] that their verifiable credentials were correctly derived from an existing eID system.

### 3.1.2 Issuance of Sensitive Claims.

**Threat 3.2.** An identity owner may not have the right to view the contents of a credential issued to them. Such as when the credential is a confidential letter of recommendation regarding them.

**Risk:** An identity owner could collect sensitive claims, violating the privacy of the issuers.

**Constraint:** Issuers should be able to issue credentials where claims are bound by a constraint that precludes parties – including the identity owner – from learning the contents of the claim.

**Mitigations:** A possible approach is *Encrypted Claims* where an issuer provides asymmetrically encrypted credentials with a secret key unknown to the identity owner. Identity owners present these credentials to service providers who know the secret key and can decrypt the claims. Guajardo et al. [67] introduced this in anonymous credentials. These constructions are limited as prior knowledge of the claim-accessing party is required but have found renewed interest through anonymous tokens [32, 42].

## 3.2 For Identity Owners

In the *trustful model*, the only threat to identity owners was the collection of credentials during transit. Here, we expand with the introduction of four new threats.

*3.2.1 Selective Disclosure. Data Minimalization,* collecting only necessary identity claims by the service providers, is a fundamental principle of SSI [5]. It is realized through the extension of *selective disclosure* at Layer 3.

**Threat 3.3.** A service provider may collect more identity claims than required, either maliciously or without harmful intentions, such as a bar collecting an identity owner's name and address during age verification.

**Risk:** Excessive collection of claims by service providers elevates the risks faced by identity owners when presenting credentials. Malicious service providers may exploit identity data for financial gain or profiling purposes [23], while even well-intentioned service providers inadvertently elevate identity loss risks with their over-collection of identity data [156].

**Constraint:** Credential presentations must allow identity owners to present a subset of the claims in their credentials where service providers may not learn anything beyond what was presented.

**Mitigations:** Selective disclosure mechanisms can be achieved either through *claim redaction* or *predicates*.

*Claim Redaction;* enables identity owners to remove claims from a verifiable credential when generating a presentation [157]. Service providers can verify the authenticity and integrity of the revealed claims without having access to the entire credential. Claim redaction describes an outcome where the service providers are only able to learn a subset of the claims that were initially included in a verifiable credential (e.g., a birthdate on a driver's license). Some approaches require that the identity owner only share the claims they want to disclose along with the accompanying proofs [47]; even though claims are not strictly being *redacted*, the outcome remains the same.

*Predicates;* allows identity owners to provide a predicate proof (TRUE/FALSE) on a claim (e.g., their birthdate claim indicates they are over 18) without revealing the actual claim value. The extent of the predicates available depends on the implementation chosen. Some implementations limit predicates to numerical ranges (e.g., age > 18) [28], while others enable more complex assertions (e.g., the listed address is within some defined city boundary) [94].

There are many proposals for implementing SSI selective disclosure categorized under three approaches: *atomic* (signing each claim individually), *hash records* (signing a hash of each claim rather than the claim itself, such that the service provider verifies the signature against the hashes of the claims and that the disclosed claim is the pre-image of the hash) [139], and *zero-knowledge proofs* (either as a custom proof designed for a specifically signature scheme [28, 141], or a generic zero-knowledge proof on the claim and credential [94, 114]). While all of these approaches enable claim redaction, only *zero-knowledge proofs* support genuine predicate selective disclosure. This functionality can be emulated by having the issuer sign a pre-computed predicate claim (i.e. *age\_over\_18 = TRUE*) during credential creation.

### 3.2.2 Unlinkability.

**Threat 3.4.** Selectively disclosed claims provide limited privacy preservation if service providers can link presented claims to the same identity owner across consecutive presentations or through collusion with other service providers. Consider a venue using presentations to track patronage or a car rental service linking claims with information from an insurance agent without consent.

**Risk:** If a presentation provided to a service provider contains an identifier unique to the identity owner, service providers can easily link an identity owner's actions [35]. Service providers may then learn more about an identity owner's identity than what the identity owner consented to.

**Constraint:** Verifiable presentations should not include any correlatable identifier the service provider can derive, such as identity owner DID or signature.

**Mitigations:** *Single-show* [14, 120] or *k-show* [93] credentials offer unlinkability only for a fixed (one or  $k$ ) number of presentations of a credential. For effective SSI integration, credentials must enable identity owners to generate an unlimited number of unlinkable presentations from the same credential. Two unlinkability levels exist: *complete* and *limited*.

*Complete Unlinkability;* To achieve complete unlinkability, where no presentations may ever be linked to the same credential, two transformations must be applied. The *removal of correlating factors* addresses the vulnerability of a correlating DID identifier in every credential. Some works suggest managing the scope of DIDs, where the identity owner creates multiple DID pseudonyms, each used for a small subset of their credentials [101]. This still enables correlation for credentials issued to the same DID. Instead, verifiable credentials should not have DIDs (or other identifiers) embedded in them; DIDs remain only as a means for DIDComm communications and encryption (Layer 2) [40] and issuer identification (Layer 1) [143]. At times, this is the extent of unlinkability afforded (Coelho et al. [38], Takemiya and Vanieiev [137]).

Still, the signature of the credential remains a correlating factor. With the *transformation of signatures*, when the identity owner is preparing a credential for presentation, they generate a new *proof* that they then provide in place of the signature from the issuer. This new proof is an unlinkable zero-knowledge proof of knowledge of the credential's signature. Here, anonymous credential signatures are used (BBS+) [24, 141], CL [26], PS [122]).

*Limited Unlinkability;* Complete unlinkability is not always desirable. In some situations, the service provider must be able to identify repeat presentations of identity owners, such as to prevent Sybil attacks [107]. If a service provider relies on SSI credentials for authentication into their service, such as through user accounts, then *limited unlinkability* is required, allowing service providers to identify repeat presentations; otherwise, a single credential could be used to create many accounts. The limited unlinkability construction from Zhang et al. [153] included a unique identifier derived from a combination of the identity owner and the service provider's identifiers. Service providers may use this persistent identifier to identify repeat presentations but can not use it to correlate presentations made to other service providers.

### 3.2.3 Key and Wallet Management.

**Threat 3.5.** A malicious party could compromise and steal credentials or keys, or the identity owner may lose their keys.

**Risk:** An identity owner who loses their keys or credentials (or has them compromised by a malicious party) then loses control over their personal identity information and its use.

**Constraint:** Credentials and keys must be protected in wallets, ensuring they are not vulnerable to being stolen or lost.

**Mitigations:** The risk of losing keys and, in turn, access to the credentials in the identity wallet are not concerns unique to SSI [106]. Potential solutions are:

*Trusted Third Party Key Recovery;* Soltani et al. [130] proposed a key recovery method using trusted third parties for SSI. The wallet of the identity owner is still in control of the user's credentials. As a precaution, third-party providers back up the keys in escrow. Approaches vary across implementations, but all either rely on some form of trusted third-party service [90, 96] or friend [96, 111].

*Multi-Factor Authentication (MFA);* With multi-factor authentication for credential presentations, the compromise of a wallet does not allow a malicious actor to present the compromised credentials as their own. Zhang et al. [153] implemented single sign-on using anonymous credentials following a structure compatible with Layer 3. Service providers could require that identity owners present the same credentials from multiple devices.

*Multiple Devices;* The same design used by Zhang et al. [153] for *multi-factor authentication* also enables *multiple devices*. Since credentials can be stored on any number of devices, identity owners could present their credentials from any registered device.

**Threat 3.6.** Identity owners cannot always be expected to follow best practices regarding safeguarding and maintaining access to their wallets. Without proper consideration, identity owners could rely on insecure wallet protection mechanisms (e.g., credentials stored on devices without access protection through passwords).

**Risk:** Sacrificing credential security and sovereignty for usability places identity owners at risk of damaging their self-sovereignty (losing control of its use) or leakage of their credential data.

**Constraint:** Identity owners must maintain control over their credentials and keys regardless of how they are stored [5].

**Mitigations:** A solution such as *Cloud Wallets;* places the entire responsibility of credential management with a cloud provider. Naively trusting a cloud provider with control over identity claims and their disclosure would violate the self-sovereignty of the identity owners (the cloud provider can dictate credential use). Furthermore, Lagutin et al. [91] identified that constrained (e.g., embedded) devices might be unable to handle the resource overhead required by verifiable credentials. Their work ensured the constrained devices maintained control over their DIDs, but the credential processing was delegated to an authorization server. Notably, this implementation requires potentially unfounded trust in the authorization server and is only appropriate for SSI deployments in limited scopes.

To address these trust concerns, Krenn et al. [88], later improved by Haböck and Krenn [68] presented anonymous credentials for the cloud. The proposed cloud wallets could never access the stored attributes in the clear, even when presenting the credentials. Their implementation utilized proxy re-encryption [11],

structure-preserving signatures [59], and zero-knowledge proofs of knowledge [11]. To provide transparency of a cloud wallet's actions, Chase et al. [33] proposed a credential transparency system that allowed an identity owner to audit the actions of their cloud wallet.

### 3.3 For Service Providers

Section 2.2 introduced several threats service providers face; here, we expand on the threats faced by service providers and introduce new threats.

**3.3.1 Non-Transferability.** Previously, as a means to address non-transferability (Threat 2.3), the concept of embedding the identity owner DID in the credential during issuance was introduced. However, this is not acceptable when *unlinkability* needs to be preserved. Here, we expand the mitigation approaches towards non-transferability through proving *what you know* or *what you own* with *linkable* and *unlinkable* variations of each.

**What You Know;** DIDs embedded in a verifiable credential provide a suitable identifier to prove possession of a credential. This approach is the simplest and most often used [109]. This is *visible non-transferability*, providing service providers with a linkable identifier in every credential presentation.

A more robust solution is required to preserve *unlinkability*; we classify these as *hidden*. Camenisch and Lysyanskaya [28] presented protocols where the identity owners commit to a link secret when requesting a credential. The issuer verifies that the identity owner knows the committed link secret without learning the secret and signs the commitment. When identity owners present their credentials, they recommit to their link secret and share this pseudonym with the service provider. The service provider checks that the identity owner owns the committed secret. The identity owner also proves that the issuer signed a commitment on the same hidden secret. BBS signatures follow a similar principle [104]. Alternatively, the *BoundBBS* proposal uses BLS signatures [19] to bind credentials through knowledge of the secret BLS key without using the blind signing feature of BBS. Using BLS signatures removes the need to manage the representation of the link secrets and simplifies the BBS specification [145].

**What You Have;** Physical ID cards often have photos of the holder to provide these same proof-of-possession assurances using biometrics. Othman and Callahan [119] proposed the Horcrux Protocol, an SSI construction based on biometrics. This construction established DIDs linked to biometric factors, tying credentials issued to a specific identity owner. However, this implementation did not provide the *unlinkability* features and is a *visible biometric* solution.

As a *hidden* implementation of biometric non-transferability, Impagliazzo and More [85] introduced a scheme that tied anonymous credentials to biometric identifiers. Their construction enabled non-transferability without compromising the privacy of the identity owner. Adams [4] and Gerdes et al. [61] presented improvements on this scheme.

Credentials may also be tied to the identity owner's specific device (i.e., the device that hosts the credential wallet). These identifiers do not always have to be public/private key pairs: in the Internet of Things (IoT) context, Niya et al. [116] exploited the manufacturing variability of devices to produce a unique device identifier to embed in credentials.

### 3.3.2 Credential Validity (Revocation).

**Threat 3.7.** Credentials have a lifespan; identity owners can lose the privileges associated with a credential (e.g., having the right to drive revoked or a university identity card expiring after a student graduates).

**Risk:** A service provider who accepts a credential that the issuer has revoked is vulnerable to accepting identity owners who no longer have authority over the credentials they present.

**Constraint:** Service providers must be able to determine if the credentials they are presented are still valid.

**Mitigations:** We can broadly categorize the mitigating approaches as either *expiring* or *immediate*.

With *Expiring* credential validity, the issuer specifies a validity period. When a service provider is presented with a credential, they reject any credentials for which the validity period has passed. Credential validity can be renewed through the re-issuance of the credentials. Rather than requiring the whole credential to be re-issued, issuers can instead generate lightweight attestations to the validity of the credential that accompanies the presentation. Work by Abraham et al. [2] proposed a revocation scheme where the Layer 1 network provides these attestations. Camenisch et al. [25] provided a similar construction, where issuers regularly update and publish values that identity owners retrieve and use to re-validate their expired credentials.

These short lifespans allow for efficient verification of credentials. However, they are unable to provide immediate revocation, and a balance must be struck between having a short enough validity period that credentials can be revoked in an acceptable time frame while not being too short that issuers are over-burdened re-issuing or attesting to credential validity.

With *Immediate* revocation, the validity of a credential is altered with direct effect. Commonly, this is accomplished through a publicly available list of credential statuses (either listing the revoked or valid credentials). Issuers update this record if a credential must be revoked; service providers reference the list when verifying a credential. Abraham et al. [3] proposed using a revocation list utilizing the existing Layer 1 distributed ledger.

Rather than maintaining a list of revoked or valid credentials, accumulators can collect the identifiers of all valid credentials into a single value that is published [36]. When a credential is revoked, the accumulator is updated to no longer include the revoked credential. To preserve the *unlinkability* requirement, the identity owner proves in zero-knowledge that their credential is registered in the dynamic accumulator [27].

Instead of a list or accumulator – but still with direct effect – *reclaimID*, from Schanzenbach et al. [124] used the online nature of their identity records to update the record's status.

Compared to expiry credential validity, immediate revocation removes the burden of re-issuing credentials or re-publishing values for the issuers. In exchange, the identity infrastructure must accommodate the revocation or status list.

### 3.3.3 Identity Owner Identification.

**Threat 3.8.** Identity owners may not always be collecting credentials from trusted issuers. There may be situations where identity owners self-published claims of their identity.



**Risk:** Service providers are provided with no indication of the trustworthiness of an identity owner’s self-published credentials.

**Constraint:** Presentations of self-published credentials should provide some confidence for service providers in the claims truth.

**Mitigations:** *Reputation of Self Published Claims* is an approach where reputation is built by collecting attestations to the validity of credentials in a web-of-trust approach [135]. Aside from variations across implementations and how reputation is collected and measured, implementations vary in their motivations as to why there is no trusted issuer.

**Threat 3.9.** Identity owners may need to delegate the authority of their credentials to another identity owner in an exceptional circumstance (e.g., in healthcare).

**Risk:** Without delegation capabilities, identity owners who are in situations where they are unable to present their credentials are essentially left without their identity.

**Constraint:** Credentials should enable their authority and claims to be shared with another identity owner when required.

**Mitigations:** For *Delegation of Credentials*, the identity owner takes on the role of an issuer, passing on the rights of their credential to a new identity owner [110]. Delegation can follow a chain, such as in the example of a hospital that issues a credential to a doctor who then issues the rights to their associates when they are on leave [43] or a patient who delegates a letter of their authority to a friend or family member [110]. Chase and Lysyanskaya [34] (and Demuynck et al. [43]) introduced delegatable anonymous credentials as a further constraint that the delegating issuer should not be identifiable from the presentation of a delegated credential.

**Threat 3.10.** The ownership of a credential issued by a *trustworthy* issuer does not indicate that the identity owner will behave appropriately. From Sections 3.2.1 (*selective disclosure*) and 3.2.2 (*unlinkability*), the service provider may not even be able to identify a misbehaving identity owner. Consider a case where SSI is used when renting a car; a driver’s license credential, while proving identity and a right to drive, offers no guarantee of safe driving.

**Risk:** Service providers may enter into a relationship with untrustworthy identity owners, and become unacceptably vulnerable.

**Constraint:** Service providers must be afforded some level of recourse or protection against identity owners who misbehave<sup>3</sup>.

**Mitigations:** Using *Blacklisting and Reputation*, when *unlinkability* is not considered, a service provider can easily blacklist a misbehaving identity owner from their service through their DID. With *unlinkability*, this becomes a more complex problem. Additionally, one service provider may wish to report their experiences to other service providers while maintaining the identity owner’s privacy. Tsang et al. [144] proposed a scheme that supported the blacklisting of anonymous credentials. Nakanishi and Kanatani [112] expanded on this by allowing service providers to contribute to a reputation score of identity owners rather than a boolean blacklist.

*Anonymity Revocation;* Another challenge introduced by *unlinkability* is how service providers can take recourse (beyond just blacklisting from their service) against misbehaving identity owners. With roots in group signatures [18], *anonymity revocation* enables a service provider, supported by a mutually trusted third

<sup>3</sup>The concept of misbehaving is situational to the identities use and the service provider’s policies.

party between the identity owner and service provider, to revoke the anonymity provided by *unlinkability*. Camenisch and Lysyanskaya [26] proposed that identity owners include some traceable identifier encrypted against the public key of a trusted revocation manager. The encrypted identifier included de-anonymization conditions to ensure the service provider could not tamper with the agreed-upon de-anonymization policy between the service provider and the identity owner.

**3.3.4 Issuer Identification.** The authority to issue credentials is not enforced through SSI; this falls on the human trust and governance layer of SSI (Layer 4) [56]. As a brief expansion, some SSI proposals provide a means for the service provider to build trust in issuers that may not be *traditionally* recognized via reputation.

*Issuer Reputation;* From the same motivation of Threat 2.2, *issuer reputation* mechanisms enable service providers to learn some calculated reputation scores about an issuer. This reputation aims to measure the *trustworthiness* of an issuer’s attestations. Bhattacharya et al. [16] used the number of credentials issued by the issuer in a time window. Additionally, reputation may be measured against the specific claims of a credential rather than the issuer [142].

**3.3.5 Offline Verification.**

**Threat 3.11.** The identity owner’s control of credentials in SSI lends itself well to offline operations where the identity owner or service provider may not have reliable Layer 1 network access (e.g., verification of a credential in a remote region).

**Risk:** Service providers may be denied from accepting (or identity owners presenting) valid credentials if their Layer 1 network connection is interrupted or unavailable.

**Constraint:** The Layer 3 credential exchange between identity owners and service providers must only rely on Layer 2 peer-to-peer communications and not require access to the Layer 1 network.

**Mitigations:** The self-contained proofs (signatures) in verifiable credentials permit offline verification. The peer-to-peer communications of Layer 2 and pair-wise DIDs ensure neither the service provider nor the identity owner is required to reference the Layer 1 network during credential exchange. What is lost is the service provider’s ability to de-reference issuer DIDs, find credential structures, and check revocation registries. Solutions to the DID and structure challenges are agnostic of implementation but introduce further constraints on the situations in which *offline verification* could be performed. If the service provider is offline, they must know the credential structure they are requesting and have de-referenced the issuer’s DID to their public key before being offline.

In many cases, this is acceptable, especially if the service provider has intermittent access to the network to update their local cache. What poses a more significant challenge is Threat 3.7 (*credential validity*). Without immediate access to a revocation or status list, *immediate* credential validity is incompatible. Credential status options are likely constrained to *expiry* [2]. Furthermore, *key and wallet management* cannot rely on possibly inaccessible *cloud wallets* (Threat 3.6). Some *biometric* implementations for *non-transferability* (Threat 2.3) rely on server access for secure verification of credentials [119], making them incompatible with *offline verification*.

## 4 THE TRUST MODELS

This section analyses the extended trust requirements, presented as threats and mitigations in Section 3, and introduces two additional models as alternatives to the *trustful* model (Section 2.2).

### 4.1 Threat Dependencies

The discussions in Sections 3.1, 3.2, and 3.3 revealed implicit dependencies between threats as they emerged. Figure 3 summarizes these dependencies.




Figure 3: SSI Threat Dependencies

The privacy protection provided by selective disclosure is compromised if a malicious or colluding service provider may correlate presentations. This necessitates the inclusion of Threat 3.4 (*unlinkability*) if Threat 3.3 (*selective disclosure*) is included in the trust assumptions.

Threat 2.3 (*non-transferability*) necessitates non-transferable credentials. Approaches that utilize an embedded identifier visible to service providers introduce a correlatable identifier that enables linking. In the presence of Threat 3.4 (*unlinkability*), *hidden* non-transferability constructions must be employed to mitigate this risk.

Similarly, for Threat 3.7 (*credential validity*), approaches that employ credential validity lists must exclude correlatable identifiers from the lists and instead adopt an approach where identity owners prove their membership in an accumulator.

The anonymity potentially provided by *unlinkability* (Threat 3.4) may place the service providers at risk, suggesting Threat 3.10 (*blacklisting, reputation, and anonymity revocation*).

The incorporation of Threat 3.11 (*offline verification*) imposes limitations on Threat 3.7 (*credential validity*), restricting implementations that require Layer 1 access (immediate via lists).

The consideration of these dependencies is relevant to discussions involving alternative models and the range of trust assumptions within each model. Opting to address one threat may require incorporating another related threat mitigation.

### 4.2 The Trust Models of SSI

The design of SSI components is influenced by underlying trust assumptions, often implicitly embedded in academic literature and industry implementations. We identify and categorize these implicit trust assumptions into three distinct trust models for SSI, representing the overarching approaches to trust. These trust models define the scope of what is untrusted in the Layer 4 application; effectively, the models identify the threats users face. The three models represent varying levels of trust, with each subsequent model assuming less trust, accepting less risk, and subsequently capturing more threats than the previous. We further explore the

Table 3: Trust Models and Recommended Threats

Extension / Feature	Threat	Trustful	Intermediate Trust	Zero-Trust
Identity Owner Identification	2.1	●	●	●
	3.8	○	○	●
	3.9	○	○	○
	3.10	○	●	●
Issuer Identification	2.2	●	●	⊗
Non-Transferability	2.3	●	●	●
Protected Communication	2.4	●	●	●
	2.5	●	●	●
Identity On-Boarding	3.1	○	○	⊗
Issuance of Sensitive Claims	3.2	○	○	○
Selective Disclosure	3.3	○	●	●
Unlinkability	3.4	○	●	●
Key and Wallet Management	3.5	○	●	●
	3.6	○	●	●
Credential Validity (Revocation)	3.7	●	●	●
Offline Verification	3.11	○	○	○

● required ● likely ○ optional extension ⊗ forbidden

implications of these assumptions and provide recommendations on what threats are captured by each model in the *Implications on Threat Mitigation* sections. Table 3 summarizes the recommended threats captured by each trust model, providing a reference of the remaining discussions in this section.

#### 4.2.1 Trustful Model.

*General Concept.* The *trustful* model (from Section 2.2) assumes as few threats as possible. Implementation approaches that follow this trust model often *do not consider the sensitivity of the identity data to be high*. This model assumes *complete trust in issuers* regarding credential issuance and indicates *no trust concerns about service providers' behaviour*.

*Implications on Threat Mitigation.* As the most trusting model, the captured threats must provide a functional identity system. Credentials should be verifiable by service providers (Threat 2.1 *identity owner identification*). Service providers should be able to ascertain the issuer of a credential (Threat 2.2 *issuer identification*) and the ownership of credentials (Threat 2.3 *non-transferability*). Additionally, the Layer 2 communication channels employed and the Layer 1 identification should provide assurances of security (Threats 2.4 and 2.5 *protected communication*).

Implementing an SSI system under this model necessitates a rudimentary Layer 1 DID infrastructure for issuer identification, Layer 2 protocols like DIDComm [40] to safeguard communications, and a credential format coupled with a *simple* verification mechanism such as an ECDSA signature [87].

#### 4.2.2 Intermediate Trust Model.

*General Concept.* The *intermediate trust model* extends upon the *trustful model* with a lower risk tolerance, making it more suitable for applications handling sensitive data. Service providers are no longer trusted, and all entities adopt a more cautious approach to potential threats. Implementations adhering to the *intermediate trust model* often target large-scale identity systems, such as those used for government-issued credentials, as opposed to internal company authentication tools.

*Implications on Threat Mitigation.* To safeguard identity owners from the now untrusted service providers, Threat 3.3 (*selective*

*disclosure*) should be considered. The lower risk tolerance also indicates that Threat 3.7 (*credential validity*) should be included. Following the dependencies outlined in Figure 3, Threat 3.4 (*unlinkability*) would also need to be included. The implementation of *credential validity* (Threat 3.7) and Threat 2.3 (*non-transferability*) must be restricted to their *unlinkable/hidden* variants, and Threat 3.10 (*black-listing, reputation, and anonymity revocation*) may be considered viable.

Implementations do not necessitate modifications to Layer 1 or 2. Layer 3 would require the selection of a signature scheme that supports *selective disclosure* and *unlinkability*. Furthermore, the format and protocols may need to be extended to accommodate the additional extensions.

#### 4.2.3 Zero-Trust Model.

*General Concept.* The *zero trust model* challenges the notion of a predefined set of trusted issuers with absolute authority to verify identities. Under this assumption, any form of central authority is considered a threat that demands mitigation. These systems embody a fully decentralized identity system.

*Implications on Threat Mitigation.* To eliminate the concept of trusted issuers, Threat 2.2 (*issuer identification*) is removed and prohibited from the model. Threat 3.1 (*identity on-boarding*) is also incompatible. Threat 3.8 (*self-attestation of claims*) is introduced to augment *identity owner identification*. Beyond these changes, the low-risk tolerance implies the same trust assumptions as the *intermediate trust model*.

In fully *trustless* systems, the centralized authority of traditional public key infrastructure (PKI) is no longer appropriate and, if present, is replaced with decentralized trust anchors, typically in the form of a verifiable data registry (VDR) such as a blockchain.

The absence of established issuers necessitates reconsidering the Layer 3 credential exchange structure. Credentials and their corresponding proof mechanisms must support alternative assurance mechanisms, such as web-of-trust and self-attestation of claims.

## 5 THE STATE OF SSI

This section examines the current landscape of self-sovereign identity (SSI) from the perspective of the components and threat assumptions discussed earlier. Both academic literature and industry implementations are evaluated. The component selections and the approaches adopted for the extensions are presented. Columns were omitted in instances where no information was available.

Table 5 presents a compilation of academic contributions gathered through the literature review outlined in Appendix A. The double horizontal line separates works that present a complete SSI construction (top) from those that only introduce a variation of an extension (bottom). Table 6 lists industry implementations. These implementations are commonly referred to as credential profiles [89]. In both tables, (–) indicates that the component or extension was not considered in the implementation. Table 4 further categorizes the approaches employed for the extensions in the credential profiles.

**Table 4: Credential Profile Extension Categorizations**

Selective Disclosure	None	RS256 [128]	
	Redact	ECDSA* [87], BBS(+) [141], BoundBBS [145], GGM-Merkle [109]	
	Predicates	CL [28], ZKP (zkSNARK [114])	
Unlinkability	None	ECDSA, RS256, GGM-Merkle, BLS	
	Limited	–	
	Complete	CL, BBS(+), BoundBBS, ZKP (zkSNARK)	
Non-Transferability	What you know	Linkable	DID Key, JWT [13], Any Key
		Hidden	Link Secret, BLS Key (BoundBBS)
	What you have	Linkable	Biometric or Device
		Hidden	–
Revocation	None	–	
	Expiry	Epoch Expiry	
	Immediate	AnonCreds Revocation [78], W3C Status List 2021 [99], EBSI Status List [50], JWT and CWT Status List [102]	

\* When used with SD-JWT-VC redaction is possible [139]. ECDSA does not offer selective disclosure with the other credential formats in Table 6.

Several notable distinctions between academic and industry SSI implementations underscore the divergence between the two domains. Most prominently, the choice of trust anchors differs significantly. Academic constructions predominantly rely on blockchains, such as the identity-specific Indy platform [81], as trust anchors. This extensive use of blockchains has led some to erroneously categorize SSI as a blockchain identity. However, this is inaccurate, as the majority of credential profiles in Table 6 utilize *traditional* public key infrastructure. This fundamental difference stems from varying trust assumptions and differing risk priorities. None of the industry implementations adheres to the *zero-trust* model, which disallows the existence of trusted authorities. Instead, all industry implementations permit some level of trusted entity, whereas many academic contributions do not and necessitate *zero-trust*.

Still, there are commonalities across domains. Both academic and industry approaches often share fundamental goals, such as agreeing on broad trust models for privacy protection (intermediate and zero-trust models) or aiming for a practical digital identity solution (trustful model). Despite these shared objectives, differing risk perspectives lead to variations in component choices within the overarching framework. Notably, these implementation differences do not indicate significant changes in the framework or architecture. Instead, adapting to address potential threats often requires simply swapping specific components within the established structure.

The trust assumptions listed for each contribution of Tables 5 and 6 reflect the authors' implicit threat assumptions, aligning with the models presented in Section 4. Comparing the implemented constraints or mitigations through the selected components and extensions against the recommended threats for each model category reveals that few works fully address the threats implied by the *intermediate* or *zero-trust* models.

These discrepancies do not imply any fault in the authors' implementations; however, they do emphasize the need for explicit threat definitions and careful consideration of the requisite extensions and components. Ultimately, achieving mutual agreement on a framework to capture SSI as well as a catalogue of threats and mitigations available can foster a more consistent approach to SSI across all domains.

**Table 5: Academic Contributions Comparison**

	Trust Anchor	Identification	Signature	Trust Assumption	Selective Disclosure		Unlinkability	Non-Transferability	Credential Validity	Issuer Identification	Key and Wallet Management	Identity On-Boarding	Issuance of Sensitive Claims	Identity Owner Identification	Offline Verification
Coelho et al. [38]	Blockchain	Public Key	-	Intermediate	×	-	Know	Either	-	-	-	-	-	-	-
Niu and Ren [115]	Ethereum	Public Key	-	Intermediate	×	-	Know	Immediate	-	TTP	-	-	-	-	×
Stokkink et al. [134]	Any PKI	Public Key	ZKP	Intermediate	×	✓	Have	Either	-	-	-	-	-	-	×
Laborde et al. [90]	FIDO PKI	FIDO	-	Trustful	-	-	Know	-	-	FIDO	-	-	-	-	×
Borse et al. [20]	Ethereum	Public Key	ZKP	Intermediate	-	✓	Know	-	-	-	-	-	-	-	×
Mukta et al. [109]	Parity	DID	GGM-Merkle	Intermediate	×	-	Know	-	-	-	-	-	-	-	-
Luecking et al. [105]	IOTA	DID	-	Zero-Trust	-	-	Know	-	-	-	-	-	-	-	-
Stokkink and Pouwelse [135]	Blockchain	Public Key	ZKP	Zero-Trust	≤	-	Know	Expiry	-	MFA	-	-	-	Self-Publish	×
Alsayed Kassem et al. [7]	Ethereum	Public Key	-	Intermediate	×	-	Know	-	-	-	-	-	-	-	×
Lagutin et al. [91]	Indy	DID	-	Trustful	-	-	Know	-	-	-	-	-	-	Delegation	×
Abraham et al. [1]	Blockchain	DID	BLS	Intermediate	×	-	Know	Immediate	-	eID	-	-	-	-	-
Lee et al. [94]	Blockchain	DID	zkSNARK	Intermediate	≤	✓	Know	-	-	-	-	-	-	-	-
Hamer et al. [69]	Indy	DID	-	Intermediate	-	✓	Have	Immediate	-	-	-	-	-	-	-
Abraham et al. [3]	Indy	DID	-	Trustful	-	-	Know	Immediate	-	-	eID	-	-	-	-
Schanzenbach et al. [124]	GNU Name System	Public Key	-	Zero-Trust	×	-	Know	Expiry	-	-	-	-	-	Self-Publish	×
Abraham et al. [2]	Blockchain	DID	BLS	Intermediate	×	-	Know	Expiry	-	-	-	-	-	-	✓
Takemiya and Vanieiev [137]	Iroha	DID	-	Intermediate	≤	-	Know	-	-	-	-	-	-	Self-Publish	×
Cho et al. [36]	Indy	DID	CL	Intermediate	-	✓	Know	Immediate	-	-	-	-	-	-	-
Gruner et al. [64]	Blockchain	-	-	Zero-Trust	-	-	-	-	-	-	-	-	-	Self-Publish	-
Launger et al. [92]	Indy	DID	-	Intermediate	-	-	-	-	Authority	-	-	-	-	-	-
Chakravarty and Deshpande [30]	Blockchain	DID	-	Trustful	-	-	-	-	-	-	-	-	-	-	-
Mukta et al. [110]	Fabric	DID	-	Intermediate	-	-	-	-	-	-	-	-	-	Delegation	-
Bhattacharya et al. [16]	Indy	DID	-	Intermediate	-	-	Know	-	Reputation	-	-	-	-	-	-
Zolotavkin et al. [155]	-	-	-	Intermediate	≤	✓	-	-	-	-	-	-	-	-	-
Niya et al. [116]	Ethereum	Public Key	-	Trustful	-	-	Have	-	-	-	-	-	-	-	-
Soltani et al. [130]	Indy	DID	-	Trustful	-	-	-	-	-	TTP	-	-	-	-	-
Othman and Callahan [119]	Blockchain	DID	-	Trustful	-	-	Have	-	-	-	-	-	-	-	×
Linklater et al. [96]	Any PKI	Public Key	-	Trustful	-	-	-	-	-	TTP + MD	-	-	-	Delegation	✓
Gruner et al. [66]	Blockchain	DID	-	Zero-Trust	-	-	-	-	-	-	-	-	-	Self-Publish	-

Selective Disclosure: ≤ = predicate × = claim redaction Unlinkability: ✓ = complete - = none Key and Wallet Management: MD = multi-device TTP = trusted third-party MFA = multi-factor authentication FIDO = key management offloaded to FIDO components [6]

**Table 6: Credential Profiles Comparison**

Credential	Trust Anchor	Identification		Credential Exchange			Trust Assumption	S.D.	Unlink.	Non-Transferability	Credential Validity	O.V.	K & W Mgmt.
		Issuer	Identity Owner	Format	Signature	Format							
Hyperledger AnonCreds [79]	Indy DLT	did:indy	Link Secret	AnonCred	JSON	CL	Intermediate	≤	✓	Link Secret	AnonCreds Revocation	~	MD / TTP / CD
JSON-LD + BoundBBS(+) [15]	CA PKI	did:web	BLS Public Key	JSON-LD	BoundBBS(+)	Intermediate	×	✓	BLS Key	W3C Status List 2021	~	-	
JSON-LD + BBS(+) [15]	CA PKI	did:web	did:key	JSON-LD	BBS(+)	Trustful	×	-	DID Key	W3C Status List 2021	~	-	
JWT-VC [103]	Bitcoin	did:ion	did:ion	JWT-VC	ECDSA	Trustful	-	-	DID Key	W3C Status List 2021	~	-	
JWT-VC [103]	Bitcoin + CA PKI	did:web	did:ion	JWT-VC	ECDSA	Trustful	-	-	DID Key	W3C Status List 2021	~	-	
NGI Atlantic for OpenID4VCs [12]	CA PKI	did:web	did:key	JWT-VC	ECDSA	Trustful	-	-	DID Key	Epoch Expiry	~	-	
NGI Atlantic for OpenID4VCs [12]	Any PKI	did:key	did:key	JWT-VC	ECDSA	Trustful	-	-	DID Key	Epoch Expiry	~	-	
Dutch Decentralized Identity [37]	CA PKI	did:web	did:jwt	JWT-VC	ECDSA	Trustful	-	-	jwt	W3C Status List 2021	~	-	
Microsoft Entra did:web [128]	CA PKI	did:web	did:web	JWT-VC	RS256	Trustful	-	-	jwt	W3C Status List 2021	~	-	
Microsoft Entra did:ion [128]	Bitcoin	did:ion	did:ion	JWT-VC	RS256	Trustful	-	-	jwt	W3C Status List 2021	~	-	
Selective Disclosure JWT [139]	Any PKI	Any Public Key	Any Public Key	SD-JWT-VC	ECDSA	Trustful	×	-	jwt	W3C Status List 2021	~	-	
ESSIF [113]	EBSI Trust Registries	did:ebssi	did:ebssi	SD-JWT-VC	ECDSA	Trustful	×	-	DID Key	EBSI Status List	~	-	
OpenID4VC High Assurance [150]	Any PKI	Any Public Key	Any Public Key	SD-JWT-VC	ECDSA	Trustful	×	-	Any Key	JWT and CWT Status List	~	-	
ISO mDL (ISO/IEC 18013-5) [147]	x.509 + CA PKI	X.509 Public Key	COSE Public Key	MDOC	ECDSA	Trustful	×	-	Biometric and/or Device	Epoch Expiry	✓	MD	
ICAO DTC [84]	Any PKI	Any Public Key	Any Public Key	ICAO DTC	ECDSA	Trustful	-	-	Biometric and/or Passport	Certificate Revocation List	~	-	
x.509 [74]	Any PKI	Any Public Key	Any Public Key	x.509	ECDSA	Trustful	-	-	jwt	Certificate Revocation List	~	-	
IRMA (Yivi Wallet) [22]	Any PKI	Any Public Key	Link Secret	IRMA XML	CL	Intermediate	≤	✓	Link Secret	CL Revocation	✓	-	

S.D.: Selective Disclosure Unlink.: Unlinkability O.V.: Offline Verification K & W Mgmt.: Key and Wallet Management Selective Disclosure: ≤ = predicate × = claim redaction Unlinkability: ✓ = complete - = none Key and Wallet Management: MD = multi-device TTP = trusted third-party CW = cloud wallets Offline Verification: ✓ = supported ~ = possible Aggregation of Signatures, Identity On-Boarding, Issuance of Sensitive Claims, Identity Owner Identification excluded from table as no credential profile considered them.

**6 OPEN WORK**

**Post-Quantum Resistance;** Existing signature schemes employed in SSI systems, as outlined in Table 5 and 6, lack quantum-safe resilience. This vulnerability extends beyond Layer 3 signatures to encompass the cryptographic underpinnings of Layers 1 and 2, potentially jeopardizing the entire cryptographic trust architecture depicted in Figure 2. While post-quantum secure solutions have emerged for public-key management in Layers 1 and 2, Layer 3 solutions remain limited. Though Jeudy et al. [86] proposed a quantum-secure anonymous credential scheme suitable for Layer 3 signatures, Dutto et al. [47] highlighted the continued challenge of achieving efficient selective disclosure (predicates or redaction).

**Interoperability of Credentials;** Without broad interoperability of SSI credentials, identity owners will again be siloed into identity systems that have chosen compatible constructions. Interoperability must be considered for each of the layers. Layer 1 constructions that use DIDs conforming to the standards of [133] are, by design, interoperable. Any entity that needs to dereference a DID can follow the specification of the defined method and collect the standardized DIDDoc. Credential and presentation incompatibilities represent the most significant obstacle to interoperability. The protocols, signatures, credential format and definitions, claims canonicalization, and encoding methods must all be compatible.

Verifiable credentials based on JSON-LD, such as the W3C VC Data Model [132], claim compatibility with JSON credentials [152],

but this has been disputed [148]. AnonCreds employs a proprietary credential format [79], deviating from W3C verifiable credential standards. The ISO 18013-5 mobile Driver's License proposal also adheres to a proprietary format [147].

SSI interoperability remains an ongoing challenge, and to date, there is no definitive consensus on what this should look like. The disparity of implementations in Tables 5 and 6 serves as evidence.

**Key Management;** Section 3.2.3 introduced the concept of multi-factor authentication (via presentations). The proposal by Zhang et al. [153] required that identity owners present their credentials multiple times from distinct devices. This is inefficient as all that is required is that identity owners are protected against a single on-device wallet compromise.

In addition, the protection of keys used by issuers when signing credentials through key rotation has been a source of contention within the standards [71, 76]. The *online* nature of resolved DIDDocs provides many methods for issuers to rotate or revoke their signing keys. With rotation, a credential issued under a key pre-rotation is considered valid if issued when the key was valid. Revocation of a signing key indicates that the key has been compromised, and existing credentials signed under the key should not be accepted. This is further complicated as assuming an issuer is malicious and holds a compromised key, the issue date embedded in the credential cannot be trusted. *AnonCreds* relies on the Layer 1 node attesting to the issuance date, further complicating issuance and verification and introducing further challenges to *offline verification* (Section 3.3.5). Protecting these keys and their lifecycles is critical to the resilience of SSI, and yet, there is a gap in addressing this challenge.

**Transparency of Service Providers;** Service providers establish trust in issuers through either pre-existing trust relationships (using DIDs to link the trusted issuer with the credential signer) or credential reputation (Threat 3.10). They build trust in identity owners through the Layer 3 credential exchange. What is missing is the *transparency of service providers*. Identity owners must hope that the service provider they interact with is honest. Often, it is up to the identity owner to evaluate the risk of disclosure.

Some discussions of digital identities, particularly those from government authorities, propose that the issuing entity accredit service providers (see Australia's proposed digital identity law [117]). However, this approach risks further siloing identities and hindering interoperability. The current state-of-the-art for protecting identity owners from malicious service providers offers two options: limiting service providers to a pre-approved list of accredited entities or relying on *selective disclosure* and *unlinkability* to safeguard identity owners from unfettered service providers. The effectiveness of these measures remains uncertain, and there is a strong likelihood that they may not be sufficient. Consider a scenario where a service provider must collect a claim that uniquely identifies the identity owner (i.e., a student number id) because the service provider does not trust the identity owner, then any attempt at *unlinkability* mitigation's is destined to fail. To address this, the discussion of Threat 3.10 introduced some mitigations (*blacklisting*, *reputation*, and *anonymity revocation*) that may enable a service provider to *find* trust in an *unlinkable* identity owner. Yet, this balance between the competing motivations of trust and risk priorities remains under-explored, with no clear solution to appease both parties.

**Unlinkability Remains Ambiguous;** Section 3.2.2 delves into the complexities surrounding *unlinkability* in SSI. Correlating *unlinkability* and *anonymity* is a typical response when presented with *unlinkability*. Anonymous credentials – one of the core technologies in many SSI constructions – even bear the name. However, this does not adequately communicate the nuances. *Anonymity* encompasses a broader scope than *unlinkability*. To provide *anonymity* would require that the SSI construction prevent tracking of identity owners through means beyond signatures in the presentations. Correlatable identifiers can be extracted from network protocols. TOR [46] may be able to address this in an *online* credential exchange; peer-to-peer credential exchanges proposed on top of layers such as NFC, Bluetooth, and QR Codes also face this challenge and will require consideration if total *anonymity* is required.

None of this matters, if the disclosed claims place identity owners in small identity, sets such that they could be profiled regardless. It remains disputed if applications should take on the burden of *unlinkability* if *anonymity* cannot be provided (evident by the number of credential profiles in Table 6 that provide *selective disclosure* without *unlinkability*). At a minimum, future contributions must clarify their position on *unlinkability*, communicating what level of *unlinkability* their construction provides to identity owners.

## 7 CONCLUSION

As digital identities become unavoidable, Self-Sovereign Identity (SSI) offers a promising direction for individual control and privacy. Through a systematic analysis of SSI's components and examining trust assumptions and threats posed by various actors, we developed three distinct trust models. These models capture the diverse levels of trust found across different SSI implementations and literature. We highlight the multifaceted nature of trust in SSI and how various stakeholders perceive and approach trust differently. Our comprehensive catalogue of SSI components and trust-related design requirements serves as a valuable resource for practitioners and researchers.

*For practitioners.* Practitioners tasked with designing an SSI deployment are provided with a framework to build their construction (Section 2.1). The trust models from Section 4 define a broad categorization and reference of the assumed risks. The catalogue of components, design requirements, and existing implementations in Sections 2.2, 3, and 5 may be referenced to select the appropriate mitigations based on the specific threats identified in the chosen trust model.

*For researchers.* Researchers can use this same framework and catalogue of components to locate their work in the field of SSI. Additionally, we support the justification of further work through our trust models and threats and identify open, unaddressed work.

## ACKNOWLEDGMENTS

The work has been supported by the Cyber Security Research Centre Limited whose activities are partially funded by the Australian Government's Cooperative Research Centres Programme. This research has been supported by an Australian Government Research Training Program (RTP) Scholarship.

## REFERENCES

- [1] Andreas Abraham, Felix Hörandner, Olamide Omolola, and Sebastian Ramacher. 2020. Privacy-Preserving eID Derivation for Self-Sovereign Identity Systems. In *Information and Communications Security (Lecture Notes in Computer Science)*, Jianying Zhou, Xiapu Luo, Qingni Shen, and Zhen Xu (Eds.). Springer International Publishing, Cham, 307–323. [https://doi.org/10.1007/978-3-030-41579-2\\_18](https://doi.org/10.1007/978-3-030-41579-2_18)
- [2] Andreas Abraham, Stefan More, Christof Rabensteiner, and Felix Hörandner. 2020. Revocable and Offline-Verifiable Self-Sovereign Identities. In *2020 IEEE 19th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*. IEEE, Guangzhou, China, 1020–1027. <https://doi.org/10.1109/TrustCom50675.2020.00136>
- [3] Andreas Abraham, Kevin Theuermann, and Emanuel Kirchengast. 2018. Qualified eID Derivation Into a Distributed Ledger Based IdM System. In *2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications/ 12th IEEE International Conference On Big Data Science And Engineering (TrustCom/BigDataSE)*. IEEE, New York, NY, USA, 1406–1412. <https://doi.org/10.1109/TrustCom/BigDataSE.2018.00195>
- [4] Carlisle Adams. 2011. Achieving Non-Transferability in Credential Systems Using Hidden Biometrics. *Security and Communication Networks* 4, 2 (February 2011), 195–206. <https://doi.org/10.1002/sec.136>
- [5] Christopher Allen. 2016. The Path to Self-Sovereign Identity. <https://www.lifewithalacrity.com/article/the-path-to-self-sovereign-identity/>
- [6] FIDO Alliance. 2023. FIDO Alliance - Open Authentication Standards More Secure than Passwords. <https://fidoalliance.org/>
- [7] Jamila Alysaid Kassem, Sarwar Sayeed, Hector Marco-Gisbert, Zeeshan Pervez, and Keshav Dahal. 2019. DNS-IDM: A Blockchain Identity Management System to Secure Personal Data Sharing in a Network. *Applied Sciences* 9, 15 (July 2019), 2953. <https://doi.org/10.3390/app9152953>
- [8] Afshan Aman, Eddie Hedges, Jace Hensley, Sam Curren, Brian Richter, Rolson Quadras, and Juan Caballero. 2022. Wallet And Credential Interactions Editor’s Draft. <https://identity.foundation/waci-didcomm/>
- [9] Archiveddocs. 2016. Securing PKI: Monitoring Public Key Infrastructure. [https://learn.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-r2-and-2012/dn786432\(v=ws.11\)](https://learn.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-r2-and-2012/dn786432(v=ws.11))
- [10] Rachel Arnold and Dave Longley. 2019. Zero-Knowledge Proofs Do Not Solve the Privacy-Trust Problem of Attribute-Based Credentials: What If Alice Is Evil? *IEEE Communications Standards Magazine* 3, 4 (December 2019), 26–31. <https://doi.org/10.1109/MCOMSTD.001.1900027>
- [11] Giuseppe Ateniese, Karyn Benson, and Susan Hohenberger. 2009. Key-Private Proxy Re-encryption. In *Topics in Cryptology – CT-RSA 2009 (Lecture Notes in Computer Science)*, Marc Fischlin (Ed.). Springer, Berlin, Heidelberg, 279–294. [https://doi.org/10.1007/978-3-642-00862-7\\_19](https://doi.org/10.1007/978-3-642-00862-7_19)
- [12] NGI Atlantic. 2022. Next Generation SSI Standards | NGI Atlantic. <https://ngiatlantic.eu/funded-experiments/next-generation-ssi-standards>
- [13] auth0.com. 2023. JWT.IO. <http://jwt.io/>
- [14] Foteini Baldimtsi and Anna Lysyanskaya. 2013. Anonymous Credentials Light. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security (CCS ’13)*. Association for Computing Machinery, New York, NY, USA, 1087–1098. <https://doi.org/10.1145/2508859.2516687>
- [15] Greg Bernstein and Manu Sporny. 2023. Jwks-Keys-Bbs. MATTR. <https://w3c.github.io/vc-di-bbs/>
- [16] Manas Pratim Bhattacharya, Pavol Zavorsky, and Sergey Butakov. 2020. Enhancing the Security and Privacy of Self-Sovereign Identities on Hyperledger Indy Blockchain. In *2020 International Symposium on Networks, Computers and Communications (ISNCC)*. IEEE, Montreal, QC, Canada, 1–7. <https://doi.org/10.1109/ISNCC49221.2020.9297357>
- [17] Stefano Bistarelli, Francesco Micheli, and Francesco Santini. 2023. A Survey on Decentralized Identifier Methods for Self Sovereign Identity. In *The Italian Conference on CyberSecurity*, Vol. 3488. CEUR, Bari, Italy, 1–15.
- [18] Dan Boneh, Xavier Boyen, and Hovav Shacham. 2004. Short Group Signatures. In *Advances in Cryptology – CRYPTO 2004*, David Hutchison, Takeo Kanade, Josef Kittler, Jon M. Kleinberg, Friedemann Mattern, John C. Mitchell, Moni Naor, Oscar Nierstrasz, C. Pandu Rangan, Bernhard Steffen, Madhu Sudan, Demetri Terzopoulos, Dough Tygar, Moshe Y. Vardi, Gerhard Weikum, and Matt Franklin (Eds.). Vol. 3152. Springer Berlin Heidelberg, Berlin, Heidelberg, 41–55. [https://doi.org/10.1007/978-3-540-28628-8\\_3](https://doi.org/10.1007/978-3-540-28628-8_3)
- [19] Dan Boneh, Ben Lynn, and Hovav Shacham. 2001. Short Signatures from the Weil Pairing. In *Advances in Cryptology – ASIACRYPT 2001 (Lecture Notes in Computer Science)*, Colin Boyd (Ed.). Springer, Berlin, Heidelberg, 514–532. [https://doi.org/10.1007/3-540-45682-1\\_30](https://doi.org/10.1007/3-540-45682-1_30)
- [20] Yogita Borse, Anushka Chawathe, Deepti Patole, and Purnima Ahirao. 2019. Anonymity: A Secure Identity Management Using Smart Contracts. In *SSRN Electronic Journal*. SSRN, Jaipur, India, 497–504. <https://doi.org/10.2139/ssrn.3352370>
- [21] David Bradbury. 2023. Unauthorized Access to Okta’s Support Case Management System: Root Cause and Remediation. <https://cms.okta.com/dev/harfiles>
- [22] Privacy by Design Foundation. 2023. What Is IRMA? · IRMA Docs. <https://irma.app/docs/what-is-irma/>
- [23] Carole Cadwalladr and Emma Graham-Harrison. 2018. Revealed: 50 Million Facebook Profiles Harvested for Cambridge Analytica in Major Data Breach. <https://www.theguardian.com/news/2018/mar/17/cambridge-analytica-facebook-influence-us-election>
- [24] Jan Camenisch, Manu Drijvers, and Anja Lehmann. 2016. Anonymous Attestation Using the Strong Diffie Hellman Assumption Revisited. In *Trust and Trustworthy Computing (Lecture Notes in Computer Science)*, Michael Franz and Panos Papadimitratos (Eds.). Springer International Publishing, Cham, 1–20. [https://doi.org/10.1007/978-3-319-45572-3\\_1](https://doi.org/10.1007/978-3-319-45572-3_1)
- [25] Jan Camenisch, Markulf Kohlweiss, and Claudio Soriente. 2010. Solving Revocation with Efficient Update of Anonymous Credentials. In *Security and Cryptography for Networks*, David Hutchison, Takeo Kanade, Josef Kittler, Jon M. Kleinberg, Friedemann Mattern, John C. Mitchell, Moni Naor, Oscar Nierstrasz, C. Pandu Rangan, Bernhard Steffen, Madhu Sudan, Demetri Terzopoulos, Doug Tygar, Moshe Y. Vardi, Gerhard Weikum, Juan A. Garay, and Roberto De Prisco (Eds.). Vol. 6280. Springer Berlin Heidelberg, Berlin, Heidelberg, 454–471. [https://doi.org/10.1007/978-3-642-15317-4\\_28](https://doi.org/10.1007/978-3-642-15317-4_28)
- [26] Jan Camenisch and Anna Lysyanskaya. 2001. An Efficient System for Non-Transferable Anonymous Credentials with Optional Anonymity Revocation. In *Advances in Cryptology – EUROCRYPT 2001 (Lecture Notes in Computer Science)*, Birgit Pfitzmann (Ed.). Springer, Berlin, Heidelberg, 93–118. [https://doi.org/10.1007/3-540-44987-6\\_7](https://doi.org/10.1007/3-540-44987-6_7)
- [27] Jan Camenisch and Anna Lysyanskaya. 2002. Dynamic Accumulators and Application to Efficient Revocation of Anonymous Credentials. In *Advances in Cryptology – CRYPTO 2002*, Gerhard Goos, Juris Hartmanis, Jan Van Leeuwen, and Moti Yung (Eds.). Vol. 2442. Springer Berlin Heidelberg, Berlin, Heidelberg, 61–76. [https://doi.org/10.1007/3-540-45708-9\\_5](https://doi.org/10.1007/3-540-45708-9_5)
- [28] Jan Camenisch and Anna Lysyanskaya. 2003. A Signature Scheme with Efficient Protocols. In *Security in Communication Networks (Lecture Notes in Computer Science)*, Stelvio Cimato, Giuseppe Persiano, and Clemente Galdi (Eds.). Springer, Berlin, Heidelberg, 268–289. [https://doi.org/10.1007/3-540-36413-7\\_20](https://doi.org/10.1007/3-540-36413-7_20)
- [29] Kim Cameron. 2005. The Laws of Identity.
- [30] Dipto Chakravarty and Tushar Deshpande. 2018. Blockchain-Enhanced Identities for Secure Interaction. In *2018 IEEE International Symposium on Technologies for Homeland Security (HST)*. IEEE, Woburn, MA, 1–4. <https://doi.org/10.1109/THS.2018.8574130>
- [31] CHAPI. 2023. CHAPI Credential Handler API. <https://chapi.io/>
- [32] Melissa Chase, F. Betül Durak, and Serge Vaudenay. 2023. Anonymous Tokens with Stronger Metadata Bit Hiding from Algebraic MACs. In *Advances in Cryptology – CRYPTO 2023 (Lecture Notes in Computer Science)*, Helena Handschuh and Anna Lysyanskaya (Eds.). Springer Nature Switzerland, Cham, 418–449. [https://doi.org/10.1007/978-3-031-38545-2\\_14](https://doi.org/10.1007/978-3-031-38545-2_14)
- [33] Melissa Chase, Georg Fuchsbaue, Esha Ghosh, and Antoine Plouviez. 2022. Credential Transparency System. In *Security and Cryptography for Networks (Lecture Notes in Computer Science)*, Clemente Galdi and Stanislaw Jarecki (Eds.). Springer International Publishing, Cham, 313–335. [https://doi.org/10.1007/978-3-031-14791-3\\_14](https://doi.org/10.1007/978-3-031-14791-3_14)
- [34] Melissa Chase and Anna Lysyanskaya. 2006. On Signatures of Knowledge. In *Advances in Cryptology – CRYPTO 2006 (Lecture Notes in Computer Science)*, Cynthia Dwork (Ed.). Springer, Berlin, Heidelberg, 78–96. [https://doi.org/10.1007/11818175\\_5](https://doi.org/10.1007/11818175_5)
- [35] Zhide Chen and Li Xu. 2006. Anti-Collusion Anonymous Credentials Scheme. In *The Sixth IEEE International Conference on Computer and Information Technology (CIT’06)*. IEEE, Seoul, 150–150. <https://doi.org/10.1109/CIT.2006.52>
- [36] Kang Woo Cho, Byeong-Gyu Jeong, and Sang Uk Shin. 2021. Verifiable Credential Proof Generation and Verification Model for Decentralized SSI-Based Credit Scoring Data. *IEICE Transactions on Information and Systems* E104.D, 11 (November 2021), 1857–1868. <https://doi.org/10.1587/transinf.2021NGP0006>
- [37] Dutch Blockchain Coalition. 2023. Dutch Decentralized Identity Profile (DDIP). DutchBlockchainCoalition. <https://github.com/DutchBlockchainCoalition/DDIP>
- [38] Pedro Coelho, Andre Zuquete, and Helder Gomes. 2018. A Propose for a Federated Ledger for Regulated Self-Sovereignty. In *2018 13th Iberian Conference on Information Systems and Technologies (CISTI)*. IEEE, Caceres, 1–4. <https://doi.org/10.23919/CISTI.2018.8399301>
- [39] Antea Cule. 2023. Louisiana Age Verification Law: Challenges and Solutions.
- [40] Sam Curren, Tobias Looker, Oliver Terbu, Kyle Den Hartog, Baha Shaaban, Drummond Reed, Steve McCown, Troy Ronda, George Aristy, Vyacheslav Gudkov, Alexander Shcherbakov, Alexander Martynov, Daniel Buchner, Devin Fisher, Ori Steele, Brian Richter, Juan Caballero, Lior Margalit, Timo Glastra, Andrew Whitehead, Nader Helmy, Markus Sabadello, Patrick McClurg, Stephen Curran, Daniel Hardman, and Moe Jangda. 2023. DIDComm Messaging Specification v2 Editor’s Draft. <https://identity.foundation/didcomm-messaging/spec/>
- [41] Ivan Damgård. 2002. On  $\Sigma$ -protocols. *Lecture Notes, University of Aarhus, Department for Computer Science* (2002), 84.
- [42] Alex Davidson, Ian Goldberg, Nick Sullivan, George Tankersley, and Filippo Valsorda. 2018. Privacy Pass: Bypassing Internet Challenges Anonymously.

- Proceedings on Privacy Enhancing Technologies* 2018, 3 (June 2018), 164–180. <https://doi.org/10.1515/popets-2018-0026>
- [43] Liesje Demuynck, Bart De Decker, and Wouter Joosen. 2007. A Credential-Based System for the Anonymous Delegation of Rights. In *New Approaches for Security, Privacy and Trust in Complex Environments*, Hein Venter, Mariki Eloff, Les Labuschagne, Jan Eloff, and Rossouw Von Solms (Eds.). Vol. 232. Springer US, Boston, MA, 169–180. [https://doi.org/10.1007/978-0-387-72367-9\\_15](https://doi.org/10.1007/978-0-387-72367-9_15)
- [44] Kyle Den Hartog. 2022. Indy DID Method. <https://hackmd.io/@kdenhartog/SieUS2BQw>
- [45] Oskar Deventer, Christian Lundkvist, Márton Csernai, Kyle Den Hartog, Markus Sabadello, Dan Gisolfi, Mike Varley, Sven Hammann, John Jordan, Lovesh Harchandani, Devin Fisher, Tobias Looker, Brent Zundel, Stephen Curran, Dennis Rybas, and Alexander Shcherbakov. 2023. Peer DID Method Specification. <https://identity.foundation/peer-did-method-spec/>
- [46] Roger Dingledine, Nick Mathewson, and Paul Syverson. 2004. Tor: The [Second-Generation] Onion Router. In *13th USENIX Security Symposium (USENIX Security 04)*. USENIX Association, San Diego, CA, 18.
- [47] Simone Dutto, Davide Margaria, Carlo Sanna, and Andrea Vesco. 2022. Toward a Post-Quantum Zero-Knowledge Verifiable Credential System for Self-Sovereign Identity.
- [48] EBSI. 2023. DID Method for Legal Entities. <https://hub.ebsi.eu/vc-framework/did/did-methods/legal-entities>
- [49] EBSI. 2023. Issuer Trust Model | EBSI Hub. <https://hub.ebsi.eu/vc-framework/trust-model/issuer-trust-model>
- [50] EBSI. 2023. Overview of EBSI's Revocation Methods | EBSI Hub. <https://hub.ebsi.eu/vc-framework/credential-status-framework/revocation-methods>
- [51] Jens Ernstberger, Jan Lauinger, Fatima Elsheimy, Liyi Zhou, Sebastian Steinhorst, Ran Canetti, Andrew Miller, Arthur Gervais, and Dawn Song. 2023. SoK: Data Sovereignty. In *2023 IEEE 8th European Symposium on Security and Privacy (EuroS&P)*. IEEE Computer Society, Delft, Netherlands, 122–143. <https://doi.org/10.1109/EuroSP57164.2023.00017>
- [52] Decentralized Identity Foundation. 2023. Decentralized-Identity/Ion. Decentralized Identity Foundation. <https://github.com/decentralized-identity/ion>
- [53] Decentralized Identity Foundation. 2023. Ethr DID Resolver. Decentralized Identity Foundation. <https://github.com/decentralized-identity/ethr-did-resolver>
- [54] IOTA Foundation. 2023. IOTA. <https://www.iota.org>
- [55] IRMA Foundation. 2023. IRMA. <https://privacybydesign.foundation/irma-en/>
- [56] Sovrin Foundation. 2018. Sovrin: A Protocol and Token for Self Sovereign Identity and Decentralized Trust. <https://sovrin.org/wp-content/uploads/Sovrin-Protocol-and-Token-White-Paper.pdf>
- [57] Sovrin Foundation. 2019. Sovrin Glossary V3. <https://sovrin.org/wp-content/uploads/Sovrin-Glossary-V3.pdf>
- [58] Josh Fruhlinger. 2020. Equifax Data Breach FAQ: What Happened, Who Was Affected, What Was the Impact? <https://www.csoonline.com/article/567833/equifax-data-breach-faq-what-happened-who-was-affected-what-was-the-impact.html>
- [59] Georg Fuchsbauer, Christian Hanser, and Daniel Slamanig. 2019. Structure-Preserving Signatures on Equivalence Classes and Constant-Size Anonymous Credentials. *Journal of Cryptology* 32, 2 (April 2019), 498–546. <https://doi.org/10.1007/s00145-018-9281-4>
- [60] Subodh Gangan. 2015. A Review of Man-in-the-Middle Attacks. <https://doi.org/10.48550/arXiv.1504.02115> arXiv:1504.02115 [cs]
- [61] John H. Gerdes, Chin-Tser Huang, and Mohamed A. Sharaf. 2017. Incorporating Biometrics into Veiled Certificates: Preventing Unauthorized Use of Anonymous Certificates. *Electronic Commerce Research* 17, 2 (June 2017), 289–316. <https://doi.org/10.1007/s10660-016-9222-y>
- [62] Christian Gribneau, Michael Prorock, Ori Steele, Oliver Terbu, Mike Xu, and Dmitri Zagidulin. 2023. Did-Web Method Specification. <https://w3c-ccg.github.io/did-method-web/>
- [63] W3C JSON-LD Working Group. 2023. JSON-LD - JSON for Linking Data. <https://json-ld.org/>
- [64] Andreas Gruner, Alexander Muhle, Tatiana Gayvoronskaya, and Christoph Meinel. 2018. A Quantifiable Trust Model for Blockchain-Based Identity Management. In *2018 IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData)*. IEEE, Halifax, NS, Canada, 1475–1482. [https://doi.org/10.1109/Cybermatics\\_2018.2018.00250](https://doi.org/10.1109/Cybermatics_2018.2018.00250)
- [65] Andreas Gruner, Alexander Muhle, Niko Lockenvitz, and Christoph Meinel. 2023. Analyzing and Comparing the Security of Self-Sovereign Identity Management Systems through Threat Modeling. *International Journal of Information Security* 22, 5 (Oct. 2023), 1231–1248. <https://doi.org/10.1007/s10207-023-00688-w>
- [66] Andreas Gruner, Alexander Muhle, and Christoph Meinel. 2019. Using Probabilistic Attribute Aggregation for Increasing Trust in Attribute Assurance. In *2019 IEEE Symposium Series on Computational Intelligence (SSCI)*. IEEE, Xiamen, China, 633–640. <https://doi.org/10.1109/SSCI44817.2019.9003094>
- [67] Jorge Guajardo, Bart Mennink, and Berry Schoenmakers. 2010. Anonymous Credential Schemes with Encrypted Attributes. In *Cryptology and Network Security*, Swee-Huay Heng, Rebecca N. Wright, and Bok-Min Goi (Eds.). Vol. 6467. Springer Berlin Heidelberg, Berlin, Heidelberg, 314–333. [https://doi.org/10.1007/978-3-642-17619-7\\_22](https://doi.org/10.1007/978-3-642-17619-7_22)
- [68] Ulrich Haböck and Stephan Krenn. 2019. Breaking and Fixing Anonymous Credentials for the Cloud. In *Cryptology and Network Security*, Yi Mu, Robert H. Deng, and Xinyi Huang (Eds.). Vol. 11829. Springer International Publishing, Cham, 249–269. [https://doi.org/10.1007/978-3-030-31578-8\\_14](https://doi.org/10.1007/978-3-030-31578-8_14)
- [69] Tom Hamer, Kerry Taylor, Kee Siong Ng, and Alwen Tiu. 2019. Private Digital Identity on Blockchain. *CEUR* 2599 (2019), 7.
- [70] Daniel Hardman. 2019. Aries-Rfcs/Concepts/0207-Credential-Fraud-Threat-Model/README.Md · Hyperledger/Aries-Rfcs. <https://github.com/hyperledger/aries-rfcs/blob/main/concepts/0207-credential-fraud-threat-model/README.md>
- [71] Daniel Hardman. 2020. Need to Clarify Revocation vs. Rotation · Issue #386 · W3c/Did-Core. <https://github.com/w3c/did-core/issues/386>
- [72] Daniel Hardman. 2020. No Paradox Here: ZKPs Deliver Savvy Trust. <https://dhh1128.github.io/zkpcrcs/trust-paradox-rebuttal.html>
- [73] Tim Hinchliffe. 2022. WEF Pushes Digital ID, CBDC, ESG & Crackdowns on 'Misinformation': Future Focus Report. <https://www.sociable.co/business/wef-digital-id-cbdc-esg-misinformation-future-focus-report/>
- [74] R Housley, W Ford, W Polk, and D Solo. 1999. Internet X.509 Public Key Infrastructure Certificate and CRL Profile. <https://www.ietf.org/rfc/rfc2459.txt>
- [75] Carly Huitema, Daniel Bacheneimer, Darrell O'Donnell, Drummond Reed, Jacques Bikoundou, Judith Fleenor, Kaliya Young, Karen Hand, Karl Kneiss, John Jordan, Lynn Bendixsen, P. A. Subrahmanyam, Sankarshan Mukhopadhyay, Scott Perry, Victor Syntez, Vikas Malhotra, and Wenjing Chu. 2021. Introduction to Trust Over IP. <https://trustoverip.org/wp-content/uploads/Introduction-to-ToIP-V2.0-2021-11-17.pdf>
- [76] Hyperledger. 2020. Issuer Key Rotation | Anoncreds Design. <https://github.com/hyperledger/indy-node/blob/main/design/anoncreds.md>
- [77] Hyperledger. 2022. Hyperledger/Aries-Rfcs. Hyperledger. <https://github.com/hyperledger/aries-rfcs>
- [78] Hyperledger. 2023. AnonCreds Revocation. Hyperledger. <https://github.com/hyperledger/anoncreds-revocation>
- [79] Hyperledger. 2023. AnonCreds Specification. Hyperledger. <https://github.com/hyperledger/anoncreds-spec>
- [80] Hyperledger. 2023. Hyperledger Fabric. <https://hyperledger-fabric.readthedocs.io/en/latest/index.html>
- [81] Hyperledger. 2023. Hyperledger Indy - Hyperledger Indy - Hyperledger Foundation. <https://wiki.hyperledger.org/display/indy>
- [82] Hyperledger. 2023. Hyperledger Iroha. <https://iroha.readthedocs.io/en/main/>
- [83] Hyperledger. 2023. Indy SDK. Hyperledger. <https://github.com/hyperledger/indy-sdk>
- [84] Subgroup of the New Technologies Working Group (NTWG) ICAO Technical Advisory Group on the Traveller Identification Programme (TAG/TRIP). 2020. Guiding Core Principles for the Development of a Digital Travel Credential (DTC). <https://www.icao.int/Security/FAL/TRIP/PublishingImages/Pages/Publications/Guiding%20Core%20Principles%20for%20the%20Development%20of%20a%20Digital%20Travel%20Credential%20-%20-%20DTC%29.PDF>
- [85] Russell Impagliazzo and Sara Miner More. 2003. Anonymous Credentials with Biometrically-Enforced Non-Transferability. In *Proceedings of the 2003 ACM Workshop on Privacy in the Electronic Society*. ACM, Washington, DC, 60–71. <https://doi.org/10.1145/1005140.1005150>
- [86] Corentin Jeudy, Adeline Roux-Langlois, and Olivier Sanders. 2022. Lattice Signature with Efficient Protocols, Application to Anonymous Credentials.
- [87] Don Johnson, Alfred Menezes, and Scott Vanstone. 2001. The Elliptic Curve Digital Signature Algorithm (ECDSA). *International Journal of Information Security* 1, 1 (01 Aug 2001), 36–63. <https://doi.org/10.1007/s102070100002>
- [88] Stephan Krenn, Thomas Lorünser, Anja Salzer, and Christoph Striecks. 2018. Towards Attribute-Based Credentials in the Cloud. In *Cryptology and Network Security*, Srdjan Capkun and Sherman S. M. Chow (Eds.). Vol. 11261. Springer International Publishing, Cham, 179–202. [https://doi.org/10.1007/978-3-030-02641-7\\_9](https://doi.org/10.1007/978-3-030-02641-7_9)
- [89] Andre Kudra, Torsten Lodderstedt, Paul Bastian, Mirko Mollik, Maaïke van Leuken, and Caspar Roelofs. 2023. Credential Comparison Matrix. <https://github.com/vcstuff/credential-profile-comparison>
- [90] Romain Laborde, Arnaud Oglaza, Samer Wazan, Francois Barrere, Abdelmalek Benzekri, David W. Chadwick, and Remi Venant. 2020. A User-Centric Identity Management Framework Based on the W3C Verifiable Credentials and the FIDO Universal Authentication Framework. In *2020 IEEE 17th Annual Consumer Communications & Networking Conference (CCNC)*. IEEE, Las Vegas, NV, USA, 1–8. <https://doi.org/10.1109/CCNC46108.2020.9045440>
- [91] Dmitriy Lagutin, Yki Kortensniemi, Nikos Fotiou, and Vasilios A. Siris. 2019. Enabling Decentralised Identifiers and Verifiable Credentials for Constrained IoT Devices Using OAuth-based Delegation. In *Proceedings 2019 Workshop on Decentralized IoT Systems and Security*. Internet Society, San Diego, CA, 6. <https://doi.org/10.14722/diss.2019.23005>
- [92] Jan Lauinger, Jens Ernstberger, Emanuel Regnath, Mohammad Hamad, and Sebastian Steinhorst. 2021. A-PoA: Anonymous Proof of Authorization for

- Decentralized Identity Management. In *2021 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*. IEEE, Sydney, Australia, 1–9. <https://doi.org/10.1109/ICBC51069.2021.9461082>
- [93] Mohamed Layouni and Hans Vangheluwe. 2007. Anonymous K-Show Credentials. In *Public Key Infrastructure*, Javier Lopez, Pierangela Samarati, and Josep L. Ferrer (Eds.), Vol. 4582. Springer Berlin Heidelberg, Berlin, Heidelberg, 181–192. [https://doi.org/10.1007/978-3-540-73408-6\\_13](https://doi.org/10.1007/978-3-540-73408-6_13)
- [94] Jeonghyuk Lee, Jaekyung Choi, Hyunok Oh, and Jihye Kim. 2021. Privacy-Preserving Identity Management System.
- [95] Shu Yun Lim, Omar Bin Musa, Bander Ali Saleh Al-Rimy, and Abdullah Almasri. 2022. Trust Models for Blockchain-Based Self-Sovereign Identity Management: A Survey and Research Directions. In *Advances in Blockchain Technology for Cyber Physical Systems*, Yassine Maleh, Lo' ai Tawalbeh, Saad Motahhir, and Abdelhakim Senhaji Hafid (Eds.). Springer International Publishing, Cham, 277–302. [https://doi.org/10.1007/978-3-030-93646-4\\_13](https://doi.org/10.1007/978-3-030-93646-4_13)
- [96] Gregory Linklater, Christian Smith, Alan Herbert, and Barry Irwin. 2018. Toward Distributed Key Management for Offline Authentication. In *Proceedings of the Annual Conference of the South African Institute of Computer Scientists and Information Technologists*. ACM, Port Elizabeth South Africa, 10–19. <https://doi.org/10.1145/3278681.3278683>
- [97] Michael Lodder and Daniel Hardman. 2023. Sovrin DID Method Specification. <https://sovrin-foundation.github.io/sovrin/spec/did-method-spec-template.html>
- [98] T Lodderstedt, K Yasuda, and T Looker. 2023. OpenID for Verifiable Credential Issuance. [https://openid.net/specs/openid-4-verifiable-credential-issuance-1\\_0.html](https://openid.net/specs/openid-4-verifiable-credential-issuance-1_0.html)
- [99] Dave Longley and Manu Sporny. 2023. Bitstring Status List v1.0. <https://www.w3.org/TR/vc-bitstring-status-list/>
- [100] Dave Longley and Manu Sporny. 2023. Verifiable Credential Data Integrity 1.0. <https://w3c.github.io/vc-data-integrity/>
- [101] Dave Longley, Dmitri Zagidulin, and Manu Sporny. 2022. The Did-Key Method v0.7. <https://w3c-ccg.github.io/did-method-key/>
- [102] Tobias Looker and Paul Bastian. 2023. JWT and CWT Status List. Internet Draft draft-looker-oauth-jwt-cwt-status-list-01. Internet Engineering Task Force.
- [103] Tobias Looker, Andrew Hughes, David Waite, Jacob Yanicak, Valerie Lanard, Daniel Godbout, Rohit Gulati, Eric Kuhn, Jeremie Miller, Oliver Terbu, and Niels Klomp. 2023. JWT VC Presentation Profile. <https://identity.foundation/jwt-vc-presentation-profile/>
- [104] T. Looker, V. Kalos, A. Whitehead, and M. Lodder. 2023. The BBS Signature Scheme. <https://identity.foundation/bbs-signature/draft-irtf-cfrg-bbs-signatures.html>
- [105] Markus Luecking, Christian Fries, Robin Lamberti, and Wilhelm Stork. 2020. Decentralized Identity and Trust Management Framework for Internet of Things. In *2020 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*. IEEE, Toronto, ON, Canada, 1–9. <https://doi.org/10.1109/ICBC48266.2020.9169411>
- [106] David Paul Maher. 1996. Crypto Backup and Key Escrow. *Commun. ACM* 39, 3 (March 1996), 48–53. <https://doi.org/10.1145/227234.227241>
- [107] Deepak Maram, Harjasleen Malvai, Fan Zhang, Nerla Jean-Louis, Alexander Frolov, Tyler Kell, Tyrone Lobban, Christine Moy, Ari Juels, and Andrew Miller. 2021. CanDID: Can-Do Decentralized Identity with Legacy Compatibility, Sybil-Resistance, and Accountability. In *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, San Francisco, CA, 1348–1366. <https://doi.org/10.1109/SP40001.2021.00038>
- [108] Alexander Mühle, Andreas Grüner, Tatiana Gayvoronskaya, and Christoph Meinel. 2018. A Survey on Essential Components of a Self-Sovereign Identity. *Computer Science Review* 30 (November 2018), 80–86. <https://doi.org/10.1016/j.cosrev.2018.10.002>
- [109] Rahma Mukta, James Martens, Hye-young Paik, Qinghua Lu, and Salil S. Kanhere. 2020. Blockchain-Based Verifiable Credential Sharing with Selective Disclosure. In *2020 IEEE 19th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*. IEEE, Guangzhou, China, 959–966. <https://doi.org/10.1109/TrustCom50675.2020.00128>
- [110] Rahma Mukta, Hye-Young Paik, Qinghua Lu, and Salil S. Kanhere. 2022. CredTrust: Credential Based Issuer Management for Trust in Self-Sovereign Identity. In *2022 IEEE International Conference on Blockchain (Blockchain)*. IEEE, Espoo, Finland, 334–339. <https://doi.org/10.1109/Blockchain55522.2022.00053>
- [111] Nitin Naik and Paul Jenkins. 2020. uPort Open-Source Identity Management System: An Assessment of Self-Sovereign Identity and User-Centric Data Platform Built on Blockchain. In *2020 IEEE International Symposium on Systems Engineering (ISSE)*. IEEE, Vienna, Austria, 1–7. <https://doi.org/10.1109/ISSE49799.2020.9272223>
- [112] Toru Nakanishi and Takeshi Kanatani. 2018. An Efficient Blacklistable Anonymous Credential System with Reputation Using Pairing-Based Accumulator. In *2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications/ 12th IEEE International Conference On Big Data Science And Engineering (TrustCom/BigDataSE)*. IEEE, New York, NY, USA, 1140–1148. <https://doi.org/10.1109/TrustCom/BigDataSE.2018.00158>
- [113] NGI. 2022. eSSIF-Lab. <https://essif-lab.eu/>
- [114] Anca Nitulescu. 2020. Zk-SNARKS: A Gentle Introduction.
- [115] Jianlin Niu and Zhiyu Ren. 2021. A Self-Sovereign Identity Management Scheme Using Smart Contracts. *MATEC Web of Conferences* 336 (2021), 08005. <https://doi.org/10.1051/mateconf/202133608005>
- [116] Sina Rafati Niya, Benjamin Jeffrey, and Burkhard Stiller. 2020. KYoT: Self-sovereign IoT Identification with a Physically Unclonable Function. In *2020 IEEE 45th Conference on Local Computer Networks (LCN)*. IEEE, Sydney, NSW, Australia, 485–490. <https://doi.org/10.1109/LCN48667.2020.9314816>
- [117] Commonwealth of Australia. 2023. Australia's Digital ID Bill. <https://www.digitalidentity.gov.au/digital-id-bill>
- [118] Singtel Optus. 2022. Optus Notifies Customers of Cyberattack Compromising Customer Information.
- [119] Asem Othman and John Callahan. 2018. The Horcrux Protocol: A Method for Decentralized Biometric-based Self-sovereign Identity. In *2018 International Joint Conference on Neural Networks (IJCNN)*. IEEE, Rio de Janeiro, Brazil, 1–7. <https://doi.org/10.1109/IJCNN.2018.8489316>
- [120] Christian Paquin and Greg Zaverucha. 2023. U-Prove Cryptographic Specification V1.1. <https://www.microsoft.com/en-us/research/wp-content/uploads/2016/02/U-Prove20Cryptographic20Specification20V1.1.pdf>
- [121] José Parra Moyano and Omri Ross. 2017. KYC Optimization Using Distributed Ledger Technology. *Business & Information Systems Engineering* 59, 6 (Dec. 2017), 411–423. <https://doi.org/10.1007/s12599-017-0504-2>
- [122] David Pointcheval and Olivier Sanders. 2017. Reassessing Security of Randomizable Signatures.
- [123] Abylay Satybaldy, Mariusz Nowostawski, and Jørgen Ellingsen. 2020. Self-Sovereign Identity Systems. In *Privacy and Identity Management. Data for Better Living: AI and Privacy: 14th IFIP WG 9.2, 9.6/11.7, 11.6/SIG 9.2.2 International Summer School, Windisch, Switzerland, August 19–23, 2019, Revised Selected Papers*, Michael Friedewald, Melek Önen, Eva Lievens, Stephan Krenn, and Samuel Fricker (Eds.). Springer International Publishing, Cham, 447–461. [https://doi.org/10.1007/978-3-030-42504-3\\_28](https://doi.org/10.1007/978-3-030-42504-3_28)
- [124] Martin Schanzenbach, Georg Bramm, and Julian Schutte. 2018. reclaimID: Secure, Self-Sovereign Identities Using Name Systems and Attribute-Based Encryption. In *2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications/ 12th IEEE International Conference On Big Data Science And Engineering (TrustCom/BigDataSE)*. IEEE, New York, NY, 946–957. <https://doi.org/10.1109/TrustCom/BigDataSE.2018.00134>
- [125] M Schanzenbach, C Grothoff, and B Fix. 2023. The GNU Name System. <https://ltd.gninet.org/ltd0001/>
- [126] Frederico Schardong and Ricardo Custódio. 2022. Self-Sovereign Identity: A Systematic Review, Mapping and Taxonomy. *Sensors* 22, 15 (January 2022), 5641. <https://doi.org/10.3390/s22155641>
- [127] Vincent Schlatt, Johannes Sedlmeir, Simon Feulner, and Nils Urbach. 2022. Designing a Framework for Digital KYC Processes Built on Blockchain-Based Self-Sovereign Identity. *Information & Management* 59, 7 (Nov. 2022), 103553. <https://doi.org/10.1016/j.im.2021.103553>
- [128] Microsoft Security. 2023. Microsoft Entra Verified ID | Microsoft Security. <https://www.microsoft.com/en-ca/security/business/identity-access/microsoft-entra-verified-id>
- [129] S.S.Y. Shim, Geetanjali Bhalla, and Vishnu Pendyala. 2005. Federated Identity Management. *Computer* 38, 12 (December 2005), 120–122. <https://doi.org/10.1109/MC.2005.408>
- [130] Reza Soltani, Uyen Trang Nguyen, and Aijun An. 2019. Practical Key Recovery Model for Self-Sovereign Identity Based Digital Wallets. In *2019 IEEE Intl Conf on Dependable, Autonomic and Secure Computing, Intl Conf on Pervasive Intelligence and Computing, Intl Conf on Cloud and Big Data Computing, Intl Conf on Cyber Science and Technology Congress (DASC/PiCom/CBDCOM/CyberSciTech)*. IEEE, Fukuoka, Japan, 320–325. <https://doi.org/10.1109/DASC/PiCom/CBDCOM/CyberSciTech.2019.00066>
- [131] Reza Soltani, Uyen Trang Nguyen, and Aijun An. 2021. A Survey of Self-Sovereign Identity Ecosystem. *Security and Communication Networks* 2021 (July 2021), e8873429. <https://doi.org/10.1155/2021/8873429>
- [132] Manu Sporny, Dave Longley, and David Chadwick. 2022. Verifiable Credentials Data Model v1.1. <https://www.w3.org/TR/vc-data-model/>
- [133] Manu Sporny, Dave Longley, Markus Sabadello, Drummond Reed, Orrie Steele, and Christopher Allen. 2022. Decentralized Identifiers (DIDs) v1.0. <https://www.w3.org/TR/did-core/>
- [134] Quinten Stokkink, Georgy Ishmaev, Dick Epema, and Johan Pouwelse. 2021. A Truly Self-Sovereign Identity System. In *2021 IEEE 46th Conference on Local Computer Networks (LCN)*. IEEE, Edmonton, AB, Canada, 1–8. <https://doi.org/10.1109/LCN52139.2021.9525011>
- [135] Quinten Stokkink and Johan Pouwelse. 2018. Deployment of a Blockchain-Based Self-Sovereign Identity. In *2018 IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData)*. IEEE, Halifax, NS, 1336–1342. [https://doi.org/10.1109/Cybermatics\\_2018.2018.00230](https://doi.org/10.1109/Cybermatics_2018.2018.00230)



- [136] P. Syverson. 1994. A Taxonomy of Replay Attacks [Cryptographic Protocols]. In *Proceedings The Computer Security Foundations Workshop VII*. IEEE, Franconia, NH, 187–191. <https://doi.org/10.1109/CSFW.1994.315935>
- [137] Makoto Takemiya and Bohdan Vanieiev. 2018. Sora Identity: Secure, Digital Identity on the Blockchain. In *2018 IEEE 42nd Annual Computer Software and Applications Conference (COMPSAC)*. IEEE, Tokyo, Japan, 582–587. <https://doi.org/10.1109/COMPSAC.2018.10299>
- [138] Parity Technologies. 2023. Blockchain Infrastructure for the Decentralised Web | Parity Technologies. <https://www.parity.io/>
- [139] Oliver Terbu and Daniel Fett. 2023. *SD-JWT-based Verifiable Credentials (SD-JWT VC)*. Internet Draft draft-terbu-0auth-sd-jwt-vc-00. Internet Engineering Task Force.
- [140] O Terbu, T Lodderstedt, K Yasuda, A Lemmon, and T Looker. 2022. OpenID Connect for Verifiable Presentations. [https://openid.net/specs/openid-4-verifiable-presentations-1\\_0.html](https://openid.net/specs/openid-4-verifiable-presentations-1_0.html)
- [141] Stefano Tessaro and Chenzhi Zhu. 2023. Revisiting BBS Signatures.
- [142] Ivonne Thomas and Christoph Meinel. 2009. Enhancing Claim-Based Identity Management by Adding a Credibility Level to the Notion of Claims. In *2009 IEEE International Conference on Services Computing*. IEEE, Bangalore, India, 243–250. <https://doi.org/10.1109/SCC.2009.66>
- [143] Andrew Tobin. 2018. Sovrin: What Goes on the Ledger? , 12 pages. <https://sovrin.org/wp-content/uploads/2017/04/What-Goes-On-The-Ledger.pdf>
- [144] Patrick P. Tsang, Man Ho Au, Apu Kapadia, and Sean W. Smith. 2008. PEREA: Towards Practical TTP-free Revocation in Anonymous Authentication. In *Proceedings of the 15th ACM Conference on Computer and Communications Security*. ACM, Alexandria Virginia USA, 333–344. <https://doi.org/10.1145/1455770.1455813>
- [145] V. Valos. 2023. Bound BBS Signatures. <https://basileioskal.github.io/bbs-bound-signatures/draft-bound-bbs-signatures.html>
- [146] Elena Wolz, Matthias Gottlieb, and Hans Pongratz. 2021. Digital Credentials in Higher Education Institutions: A Literature Review. In *Innovation Through Information Systems (Lecture Notes in Information Systems and Organisation)*, Frederik Ahlemann, Reinhard Schütte, and Stefan Stieglitz (Eds.). Springer International Publishing, Cham, 125–140. [https://doi.org/10.1007/978-3-030-86800-0\\_9](https://doi.org/10.1007/978-3-030-86800-0_9)
- [147] Lucy Yang and Kaliya Young. 2023. Where Can the W3C VCs Meet the ISO 18013-5 mDL? <https://medium.com/@identitywoman-in-business/where-can-the-w3c-vc-meet-the-iso-18013-5-mdl-b2d450bb19f8>
- [148] Kristina Yasuda. 2023. Kristina Yasuda | LinkedIn. <https://www.linkedin.com/feed/update/urn:li:activity:7089289609011990528/>
- [149] K Yasuda, M Jones, and T Lodderstedt. 2023. Self-Issued OpenID Provider V2. [https://openid.net/specs/openid-connect-self-issued-v2-1\\_0.html](https://openid.net/specs/openid-connect-self-issued-v2-1_0.html)
- [150] K Yasuda and T Lodderstedt. 2023. OpenID4VC High Assurance Interoperability Profile with SD-JWT VC. <https://vcstuff.github.io/oid4vc-haip-sd-jwt-vc/draft-oid4vc-haip-sd-jwt-vc.html>
- [151] Kaliya Young. 2021. Verifiable Credentials Flavors Explained. <https://www.lfph.io/wp-content/uploads/2021/02/Verifiable-Credentials-Flavors-Explained.pdf>
- [152] Kaliya Young. 2023. Misinformation Stops Here: W3C VC 2.0 Supports JSON. <https://identitywoman.net/misinformation-stops-here-w3c-vc-2-0-supports-json/>
- [153] Zhiyi Zhang, Michał Król, Alberto Sonnino, Lixia Zhang, and Etienne Rivière. 2021. EL PASSO: Efficient and Lightweight Privacy-preserving Single Sign On. *Proceedings on Privacy Enhancing Technologies* 2021, 2 (April 2021), 70–87. <https://doi.org/10.2478/popets-2021-0018>
- [154] Xiaoyang Zhu and Youakim Badr. 2018. Identity Management Systems for the Internet of Things: A Survey Towards Blockchain Solutions. *Sensors* 18, 12 (December 2018), 4215. <https://doi.org/10.3390/s18124215>
- [155] Yevhen Zolotavkin, Jongkil Jay Jeong, Veronika Kuchta, Maksym Slavnenko, and Robin Doss. 2022. Improving Unlinkability of Attribute-based Authentication through Game Theory. *ACM Transactions on Privacy and Security* 25, 2 (March 2022), 12:1–12:36. <https://doi.org/10.1145/3501260>
- [156] Yixin Zou and Florian Schaub. 2018. Concern But No Action: Consumers’ Reactions to the Equifax Data Breach. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems (CHI EA ’18)*. Association for Computing Machinery, New York, NY, USA, 1–6. <https://doi.org/10.1145/3170427.3188510>
- [157] Brent Zundel. 2021. Why the Verifiable Credentials Community Should Converge on BBS+. <https://www.evernym.com/blog/bbs-verifiable-credentials/>

## A METHODOLOGY

We started our systematization with the seed article *The Path to Self-Sovereign Identity* by Allen [5]. References to this article were collected through ResearchRabbit<sup>4</sup>. These articles were filtered and included if they met the following requirements:

<sup>4</sup><https://researchrabbitapp.com/>

- (1) The work proposed an SSI implementation.
- (2) or; the work proposed a feature for an existing SSI implementation.
- (3) or; the work met the prior two requirements for some close adjacent field.

We also defined the following hard exclusion requirements:

- (1) The work did not propose any SSI features or implementations.
- (2) or; the work discussed a use case of SSI, that is, an implementation of existing SSI tooling for a specific scenario.
- (3) or; the work was purely focused on cryptographic protocols without practical application.
- (4) or; on point 3 from the inclusion requirements, the adjacent implementation did not have a clear analog to SSI (federations, PKI, authentication without attributes).

We collected all of the selected papers’ references and citations from the same database and repeated the filtering process. We continued these steps until the filtering converged.

This systematization resulted in two broad categories of SSI work: academic literature and industry implementations. We included extra consideration for the latter with additional sources of industry credential profiles<sup>5</sup> as they have received more attention in development and deployment in recent years than purely academic discussions. Losing their contributions would not provide a complete image of the state-of-the-art.

<sup>5</sup><https://github.com/WebOfTrustInfo/rwot11-the-hague/blob/master/draft-documents/credential-profile-comparison.md>