Spectrum of Engineering Sciences
ISSN (e) 3007-3138 (p) 3007-312X
Volume 4, Issue 06, 2026

ATTESTIFY: A HYBRID BLOCKCHAIN-IPFS FRAMEWORK FOR TRUSTLESS ACADEMIC CREDENTIAL VERIFICATION USING SOULBOUND TOKENS

Zain Ul Abidin[3], Muhammad Saad Feroz[2], Faizan Saleem[3]

[1,2]Department of Computer Science, University of Chakwal, Chakwal, Pakistan, 48800

3Department of Computer Science, University of Lahore, Sargodha Campus, Sargodha, Pakistan, 40100

guesswhozayn@gmail.com1 ferozsaad297@gmail.com2 faizan.saleem@cs.uol.edu.pk3

DOI: https://doi.org/10.5281/zenodo.20655578

Keywords
Blockchain, Soulbound Tokens, IPFS, Academic Credentials, Decentralised Verification, Ethereum, Verifiable Credentials, Self-Sovereign Identity, Literature Review

Article History
Received on 22 March 2026
Accepted on 01 June 2026
Published on 11 June 2026

Faizan Saleem
Department of Computer Science, University of Lahore, Sargodha Campus, Sargodha, Pakistan, 40100
faizan.saleem@cs.uol.edu.pk

Abstract

An ongoing institutional weakness is the widespread use of fraudulent credentials. Verification workflows that require institutions to communicate synchronously, through proprietary portals and unverified confirmation channels exacerbate the issue by providing no cryptographic guarantees of document integrity. To this end, research into blockchain-based credentialing has increased in significant amounts, but specific shortcomings have been evident throughout the literature: most proposed solutions address only parts of the credential lifecycle, none of them rely on non-transferable credentials for tokenisation of the credential holder, and participation in W3C interoperability standards is still in its infancy, if it exists at all. In this paper, we introduce Attestify, a hybrid decentralised system built using Ethereum smart contracts (Solidity, EVM-targeted), the InterPlanetary File System (IPFS) and non-transferable Soulbound Tokens (SBTs). Attestify provides a single open source platform for the entire credential lifecycle, issuing, verifying, revoking, and cryptographic identity binding. A structured literature review places the system within its architectural context, including Blockcerts, EduCTX, Ethereum-native architectures (CredChain, BlockMEDC, SuperCert), privacy-preserving architectures (ZKBAR-V), and permissioned architectures (Hyperledger Indy/Aries). Attestify is the only one that combines all these requirements in one system: non-transferability (for SBT), dual-anchor document integrity (SHA-256 digest and IPFS CID), batch issuance and on-chain revocation. We test standards readiness with the W3C Verifiable Credentials Data Model v2.0 and Decentralised Identifiers v1.0 standard, and set out a migration path which requires no smart contract changes. The paper then explores four open challenges: privacy via zero-knowledge proofs, cost reduction with migration to Layer-2 rollups, governance decentralisation, and convergence of standards.

1. Introduction

1.1 The Credential Verification Problem

Academic credentials, as a trust instrument, are used primarily by universities, employers, and regulatory agencies. But the way these credentials are checked is neither cryptographically secure, nor efficient in operation. A verifying party calls the issuing institution, waits a few days, or weeks, and eventually receives a confirmation from the issuing institution which cannot be verified. This results in a vicious cycle: no way to verify the channel itself, and the credibility of the credential depends on the credibility of the channel.

There is no place in the world that credential fraud does not happen. It is a phenomenon recorded by Grech and Camilleri [1] of the Joint Research Centre of the European Commission, and found to be trans-border, and occur across all economic development levels in institutions. Among the various obstacles to a successful centralised registrar system, Gräther et al. [2] highlighted single points of failure, natural scalability limits, and the asymmetric information flows between the various actors in the credential issuance-consumption-custodial process. Bhaskar et al. [3] made a systematic review of the problems and found that distributed ledger technology is the current research direction for these problems. They also pointed out, but not in a judgmental manner, that a significant gap exists between what is written about in academic research papers and the production matureness of the systems.

1.2 Blockchain-Based Credentialing: Capabilities and Constraints

At the bottom of the blockchain, its verification mechanism is reconfigured. The verifier asks a question, instead of it going through an intermediary institution, it is asking a question of an append-only hash-linked ledger that is replicated on the P2P network [4]. This combination of the ledger, content-addressed storage (IPFS [5]) and non-transferable identity tokens [6] provides a potential avenue for credential sovereignty wherein the owner of the credential controls and presents it.

Taking that theoretical model into a working infrastructure faces well known challenges. The SSTORE opcode is expensive enough on Ethereum (20,000 gas per new write on Ethereum, for a new 256-bit slot) to make full-document on-chain storage economically unviable [7]. The transparency of a public blockchain reveals all transactions and the state of the smart contracts to all participants in the network, introducing direct conflicts with student privacy laws, e.g., FERPA [8] and GDPR [9]. Non-technical administrators face key management challenges that they may not be able to endure [10]. There has been a significant amount of progress with the W3C Verifiable Credentials Working Group [11] on credential standardisation, but adoption of the standards in the wild is still far off, and interoperability between platforms is still hindered.

These obstacles have splintered the research landscape. Most work on one aspect of a set of problems, and very few attempt to solve more than one problem at a time. Three recent surveys [3, 12, 13] concluded that the field requires the use of integrated frameworks that manage integrity, privacy, scalability, usability and standards compliance in an integrated design. That need is what provides the motivation for the present work.

1.3 Research Questions

The following questions will be used to guide this study:

RQ1. What are the architectural approaches, functional characteristics and unmet needs in blockchain-based academic credentialing today?

RQ2. How does Attestify's features of non-transferability enforced by SBT, dual-anchor integrity and hybrid on/off-chain storage differ from and build upon previous research?

RQ3. What are the open problems and promising research directions at the cutting edge of this area as far as privacy, transaction cost, governance, and standards convergence are concerned?

1.4 Contributions

Four contributions have been reported. A systematic review of literature on blockchain credentialing is conducted, from initial hash anchoring solutions to the latest designs for privacy-preserving and Layer-2 solutions. Attestify is introduced as a reference implementation that brings together features which are not offered by any of the previous systems. A comparative evaluation of eight existing systems is presented to explain the state of the art, and to identify specific differentiators. Assessment of standards readiness for conformance to the W3C VC Data Model v2.0 is made and a research agenda for the field is presented.

1.5 Paper Organisation

The research methodology is described in Section II. The literature survey in Section III is done by architectural paradigm. A taxonomic classification is provided in Section IV. The Attestify framework is discussed in Section V. The comparative evaluation is reported in Section VI. The issue of standards alignment to the W3C is explored in Section VII. Findings and open challenges are discussed in Section VIII. Section IX concludes.

2. Research Methodology

The methodology used is Design Science Research (DSR) presented by Hevner et al. [14] and extended by Peffers et al. [15] . DSR organizes inquiry by six stages: problem identification, objective definition, design, development, demonstration, evaluation, and communication.

Structured queries were used to collect literature from IEEE Xplore, ACM Digital Library, Scopus, Web of Science and Google Scholar. The searched words were “blockchain AND (credential OR certificate OR diploma) AND (verification OR authentication OR issuance)”, and the search time was limited to the period from 2017 to 2025. There were 284 initial results to this query. The corpus was reduced to 47 papers following a series of title screening, abstract screening, and full-text screening. Only those systems that have explicit prototype or implementation was included; those that are only conceptually proposed were not included.

Comparative evaluation is conducted in a qualitative, structured approach for 10 dimensions garnered from the literature of credentialing stakeholders. There was no alternative system deployed independently or tested or scientifically measured; all comparisons are based on published descriptions and self-reported performance data. This reliance on secondary sources is a known limitation, which limits the extent to which any performance-related claims can be made.

Attestify itself has been developed incrementally over six months: developed a smart contract that runs on the EVM on Solidity (with solc ^0.8.20) and two client sides: one in Node.js/Express and one in React. The deployment was made on the Sepolia testnet of the Ethereum Proof-of-Stake chain to validate the functionality and test the gas costs.

3. Literature Review

Six thematic categories organise this review. Each corresponds to a distinct architectural direction or an unresolved problem that continues to shape the field.

3.1 Foundational Systems and Early Architectures

Early blockchain credentialing systems followed a common thread: Record information about the issuance of a credential (cryptographic evidence) on a distributed ledger and publish the information available for external verification. The most visible one is Blockcerts, a creation of the MIT Media Lab and Learning Machine [16]. It compiles a Merkle tree of certificate hashes and stores the Merkle root to Bitcoin in one operation (OP_RETURN), and returns the individual certificates as JSON files in accordance with the Open Badges v2 definition. The system was proven to be possible for decentralised credential verification and was also able to reach significant standards alignment for the time. Does not store hashes for each credential on-chain. An issuer-hosted list is used to handle revocation, removing the centralised trust dependency the purpose of the ledger was meant to eliminate. The SBT concept was yet to come into existence.

Sharples and Domingue [17] suggested an educational record system based on bitcoins with conceptual importance. They were the first to describe the conflict between learner control and institutional control over credentials, which would continue with each successive wave of this literature.

EduCTX [18] by Turkanović et al., which is published in IEEE Access, works on a different level of abstraction. It builds on the delegated proof-of-stake (DPoS) blockchain of Ark, with custom transactions encoding the ECTS credit awards, and emphasizing cross-institutional credit portability. The design does not support document-level integrity verification and content-addressed storage.

This generation has done what they did: ensured tamper evidence for academic records on the ledger. It did leave unspoken the issue of document integrity verification, complete lifecycle management and any binding between a credential and the identity of its holder.

3.2 Ethereum-Based Smart Contract Systems

Ethereum altered the architectural equation. Its Turing-complete virtual machine (the EVM) enables the logic for issuing, controlling and verifying access to its contracts to be encoded as on-chain smart contracts, which is qualitatively superior to the hash anchoring of Bitcoin-era schemes.

The main functionality of CredChain by Ammad-Ud-Din et al. [19] is to store certificate metadata on IPFS and to record references on Ethereum. It showed a successful integration of smart contracts and content-addressed storage for credential workflows. A significant limitation however is the fact that the ERC-721 standard is being used for the representation of credentials, which are transferable tokens by default. Academic credentials represent individually bound achievements that are not transferable and can therefore not circulate on secondary markets.

Alam et al. [20] proposed a conceptual design of combining IPFS-encrypted storage with Ethereum contracts to reduce fraud. It was published in MDPI Information, it was not a functional deployment and it did not involve threat modelling, or the SBT paradigm.

In their survey of blockchain use in higher education, Arenas and Fernandez [21] found a common thread, that most of the projects for verifying credentials remain at the proof-of-concept stage. User experience, governance, and regulatory alignment are all neglected, making the institutional deployability challenging.

The addition of fraudulent credential detection via machine learning classifiers, on top of on-chain hash anchoring, is introduced by SuperCert, a paper published by Khan et al. in IEEE Access (2024). That is a unique contribution. The on-chain revocation mechanism, identity tied tokens, and batch issuance are not included in the system.

3.3 Layer-2 and Rollup-Based Approaches

The hard economic limit of gas in Ethereum Layer-1. The volume of calldata used in a credential issuance transaction and whether it involves using storage operations can vary from 100,000 to 300,000 gas units. This linear increase in cost is not sustainable at the institutional level. This has prompted researchers to develop Layer-2 scaling solutions.

In IEEE Access (2024), BlockMEDC is deployed on an L2 zk-Rollup running on Ethereum and is aimed at certificates in higher education in Morocco. It is among the first peer reviewed credentialing system that has a working L2 deployment, and it is able to report substantial gas savings from L1 baselines. SBTs and batch issuance are not part of its design, and its validation covers only one national context.

The Dencun upgrade in March 2024 (EIP-4844) [24] brings blob-carrying transactions, also referred to as proto-danksharding, and reduces the cost of posting rollup data by 50 to 90 percent. Simply Staking [25] found that post-Dencun L2 transaction fees are approximately 6% of L1 equivalents. Optimistic Rollups (Arbitrum One, Base, Optimism) are the most through and mature type of rollups due to their fraud-proof mechanisms, whereas validity rollups (zkSync Era, Linea, StarkNet) provide more privacy with zk-SNARK or zk-STARK proofs submitted to L1.

The barriers of cost to blockchain credentialing are declining. How to offer these credential-specific primitives, SBTs, batch operations, revocation logic, will still be largely unanswered while combining the L2 infrastructure with these kinds of primitives.

3.4 Privacy-Preserving Credential Verification

Privacy is the most difficult unmet constraint. By design, public blockchains are transparent - all transactions inputs, all contract state and all events emitted by contracts are visible to all participants. This transparency allows for permissionless verification while also putting credential metadata in permanent and global display. The task of reconciling open verifiability with data confidentiality is an active and challenging area of research.

ZKBAR-V [26] on zkEVM, an arithmetic circuit that is zero-knowledge and supports selective disclosure directly addresses this problem. A verifier is someone who validates a credential without accessing it. The decrease in gas consumption is around 94% compared to L1, which is due to the amortised cost of the rollup's validity proof. These results are significant. There are no SBTs or content-addressed off chain storage, and only a single verification predicate, limiting their use to other use cases.

Mühle et al. [27] looked at the various self-sovereign identity (SSI) frameworks and found that there existed a convergence between distributed ledgers, decentralised identifiers and privacy-preserving credential protocols. They presented an idea they called “disclosure by design,” in which a holder builds a proof about an assertion without revealing the credential data used in the proof. That statement succinctly describes the desired architecture for credentialing that is private.

Sedlmeir et al. [28] took the argument one step further. They argued that privacy-first verification is not a choice, but a requirement for any system that processes personal educational information under the GDPR or other privacy regulations. Such systems would be legally hindered in the European Union until a mechanism of selective disclosure is put in place or equivalent.

3.5 Soulbound Tokens and Non-Transferable Identity

The non transferability of the identity in the case of tokens is formalized by Soulbound tokens. This idea was coined in 2022 by Weyl et al. [6]. An SBT is an ERC-721 token that is locked to a “Soul” (a blockchain account of an individual or organisation) such that it cannot be transferred to another account using the normal ERC-721 transferFrom and safeTransferFrom functions. While traditional non-fungible tokens are freely traded on open markets, SBTs are intended to carry commitments, credentials and affiliations that are cryptographically tied to the owner.

This idea was formalized at the protocol level in ERC-5192 [29]. It introduces a new locked() query function and Locked/Unlocked events to the ERC-721 interface, allowing wallets and indexing services to easily detect permanently non-transferable tokens.

There is interest in credential applications of SBTs. Park et al. [30] applied SBTs to the medical credential authentication in a privacy-aware architecture. Finally, Li et al. [31] investigated structural aspects for decentralised identity, suggesting that the trust model should be

understood in a different light when thinking of non-transferability. If the credential cannot be purchased, sold, or rechanged without the credential holder, then the system mandates an intuitive property for academic credentials that is more difficult to ensure with transferable properties.

However, SBTs have seldom been adopted within a full credentialing lifecycle system. The token layer is usually not used to handle the batch issuance, hash-based verification, and on-chain revocation. That gap of integration is a definite literature gap.

3.6 Alternative Ecosystems: Hyperledger and Permissioned Approaches

Public permissionless blockchains are not the only substrate for the credentialing research. Hyperledger Indy and Aries [32] provide a self-sovereign identity framework that is purpose-built to support native W3C Decentralised Identifiers, verifiable credential issuance and presentation, and selective disclosure using CL-signatures (Camenisch-Lysyanskaya). The key architectural distinction is that Indy is a permissioned ledger that uses a set of known validators that have been approved by the community. It has significantly different trust properties, throughput characteristics, and privacy properties compared to Ethereum's post-Merge proof-of-stake consensus (Gasper).

Preukschat and Reed [33] contended that permissioned ledgers are appropriate for institutional groups where identifiers and governance are better expressed in an institutional agreement than cryptoeconomic incentives.

Permissioned vs. permissionless - trade-offs are very real. Permissioned architectures provide the benefit of faster throughput (no gas market competition), lower latency (sub-second finality with honest supermajorities) and privacy (native selective disclosure via ZKP). They require in return that verifiers either become a member of the consortium, or trust a consortium maintained verifier. Openness and unconstrained verification are given up. When it comes to credentialing, that's a sacrifice that's weighted by an employer or admissions office that should always check a credential without prior arrangement. The study is based on public blockchain solutions, but also considers permissioned solutions as an important complementary concept.

4. Taxonomic Classification of Existing Approaches

The literature reviewed above is organised into a four-dimensional classification framework.

4.1 Classification Dimensions

Dimension 1: Ledger Type. Depending on the type of permission, systems choose between public permissionless ledgers (Bitcoin, Ethereum, Polygon, Ark) and permissioned consortium ledgers (Hyperledger Indy). Permissioned ledgers can have better throughput with adjustable privacy controls and restrict validator and verifier sets, while public ledgers provide open verifiability and censorship resistance, at the cost of increased per-transaction fees.

Dimension 2: Storage Architecture. On one end of the spectrum, all the data is stored on-chain (EduCTX encodes credit directly in transactions). Other, hybrid designs, commit fixed-size cryptographic digests to the ledger and off-chain store variable-size credential documents to IPFS content-addressed storage. Attestify, CredChain, and BlockMEDC sit in this middle-ground.

Dimension 3: Identity Binding. There are some that do not create any on-chain token (Blockcerts uses off-chain JSON; EduCTX uses custom transactions). These are NFTs that are issued by CredChain and can be transferred. The SBTs are non-transferable and awarded by Attestify alone, amongst the systems studied. The encasing strength determines whether a credential can be detached from the holder who has achieved it.

Dimension 4: Privacy Model. If transparency is not given, it is assumed to be the default. A minority uses selective disclosure using zero-knowledge proofs (ZKBAR-V with zk-SNARKs, Hyperledger Indy with CL-signatures). Alam et al.

use symmetric encryption for off-chain storage. Regulatory compliance posture and institutional willingness to adopt is driven by privacy capability.

4.2 Positioning Attestify

In this taxonomy, Attestify is located at one particular corner: public permissionless ledger, hybrid storage, SBT-based identity binding, and transparent (non-private) credential metadata. The only other system surveyed that can compare to the SBTs comes with full lifecycle management. A limitation of no privacy protecting verification cannot be overlooked. However, the hybrid architecture is modular: The ZK-SNARK circuits that run on off-chain credential data could be added on top of the on-chain contract logic without having to reengineer it.

5. The Attestify Framework

Architectural rationale is the emphasis of this section. The real value of Attestify's is the mixing of features that previous systems have dealt with separately, and the implementation details only comes up where they help to make a point.

5.1 Architecture Overview

The system is structured in three levels of trust with explicit trust boundaries. The client layer is a DApp built with React which offers user interaction and Ethereum wallet integration via MetaMask. The middleware layer, a Node.js/Express backend, orchestrates the issuance pipeline: PDF generation, SHA-256 digest computation, IPFS upload through the Pinata pinning service, Ethereum transaction assembly and submission, MongoDB persistence, and email notification. The consensus layer is the Ethereum blockchain (Sepolia PoS testnet) and the IPFS distributed storage network.

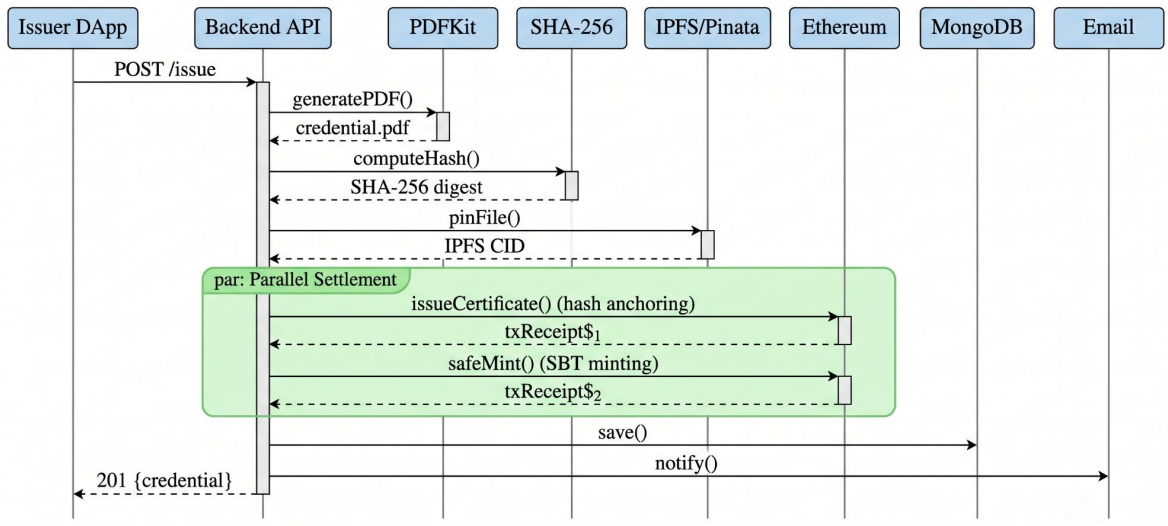

Credential issuance sequence diagram. The backend orchestrates PDF generation, SHA-256 hashing, IPFS pinning, and two parallel Ethereum transactions (hash anchoring and SBT minting) before persisting metadata and notifying the recipient.

5.2 Dual-Anchor Integrity Model

To attest integrity of the anchors, integrity is documented by two on-chain values per credential: a sha256 digest of the credential file (stored as bytes32 value) and the IPFS Content Identifier (CID) under which the file is pinned. The two anchors are a conscious design decision.

Both values are taken from the same source document in a deterministic manner. The guarantees that they offer are not independent, but redundant. This redundancy is useful in a particular way. It provides protection against implementation-level flaws in both hashing pipelines: one is the application-level SHA-256 computation that is performed by the crypto module implemented in the Node.js run-time environment; the other is the computation performed by the IPFS protocol itself (which is also SHA-256, but wrapped in a different serialisation envelope). The security of cryptographic security is not compounded beyond the second pre-image resistance provided by taking a single invocation of SHA-256.

The system will rehash the presented document with SHA-256 and compare it to the stored value on-chain, to confirm the credential. The match, along with the non-revoked status, is proof of authenticity. Formally, let d be the document presented, h be the digest of the document written into the contract storage at issuance time, and r be the boolean flag revoked (isRevoked). The check predicate is:

$$V(d) = (\text{SHA256}(d) :$$

Authenticity is true when and only when the evaluation of V(D) to true is true. It must be identical in bytes with the version anchored at issuance and there shall be no revoking on record. This predicate is deliberately imprecise. The two basic trust assumptions are the collision resistance of SHA-256 hash, and the integrity of the Ethereum state trie.

Verification is a SOLIDITY view function. It runs on the node making the request and doesn't require any Ether, gas, or a transaction, and is available to anyone with no Ethereum account..

5.3 Soulbound Token Integration

Every credential issued results in the creation of a new non-transferable ERC-721 token for the credential recipient's Ethereum wallet address. For the purposes of transferring prevention, the contract-level _update hook is overloaded with the internal OpenZeppelin's ERC-721 implementation. The “hook” is reverted whenever both the “from” and the “to” are non-zero addresses. Minting (from = address(0)) and burning (to = address(0)) pass.

This directly addresses the weakness of CredChain [19] regarding the transferability. Transferable credential tokens can be traded in NFT marketplaces, which would provide a second market for credentials. Such a prospect is just as unrealistic as what academic qualifications offer. With Soulbound enforcement, the attack vector is eliminated as ownership reassignment after minting is technically impossible.

Attestify's implementation was created prior to the widespread adoption of ERC-5192 [29] and has adopted a non-standardised approach via the OpenZeppelin v5 [38] hook mechanism. As a result, moving to ERC-5192 would enhance compatibility with wallets and indexing services that support ERC-5192. Only a small amount of code is required for the migration: a locked() query function and the emission of the proper events. The base transfer prevention logic remains the same.

5.4 Credential Lifecycle Management

The one area where Attestify is the obvious differentiator is in its breadth of life cycle. The seven stages the system implements are listed in Table 1.

Table 1: Credential Lifecycle Stages in Attestify

Lifecycle Stage	Mechanism	Description

Issuance (single)	Smart contract + IPFS	SHA-256 anchoring, CID storage, SBT minting
Issuance (batch)	Iterative contract invocations	CSV-driven batch issuance via unbounded loop
Verification (hash)	Read-only view call	Re-compute digest; compare against on-chain record
Verification (file)	File upload + re-hash	Upload document, compute SHA-256, verify on-chain
Revocation	Contract state mutation	On-chain isRevoked flag set to true (irreversible)
SBT burning	ERC-721 burn	Token destruction following revocation
Notification	Off-chain email service	Automated notification to credential recipient

This lifecycle is supported by most of the current systems, but only in parts. Blockcerts [16] does not support revocation, relying on an off-chain issuer list. EduCTX [18] deals with the transfer of credit only. CredChain [19] does not have a revocation feature for the issuance and verification of transferable tokens. There is no system which has been surveyed that incorporates all seven stages in one.

5.5 Consortium Access Control

A two-tier system of access reflects actual accreditation systems. An authorisedIssuers mapping is maintained on-chain by the contract owner, representing an accreditation authority, to keep a registry of all the authorized issuers. State-changing operations can only be performed by addresses in that registry, or by the owner of that address. Verification, on the other hand, is an unencrypted view call that does not require authentication.

Currently the administrative control is in the hands of one externally owned account (EOA). That is one place that might be compromised. If the owner's private key is exfiltrated, an attacker is able to have full control over the management of the issuer. Multi-institutional production deployments would better be implemented using multi-signature governance (Gnosis Safe) or DAO-based collective authorisation. Section VIII discusses this limitation.

5.6 Server-Side Transaction Relay

There are well-known usability issues with direct interaction with the blockchain [10]. They are dealt with using a server side relay: there is an EOA in the backend with a private key stored on the server.

The blockchain is never directly interacted with by any institutional users. They use a standard web interface and the backend builds, signs and pushes all transactions for them.

This relay brings in a trusted third party. The institutional user assumes that the back end operator will not change the transaction parameters. It is an empirical question whether that trust assumption is warranted by the usability benefit. There is no user study that addresses the question. The general issue of decentralisation and accessibility is one that is discussed throughout the credentialing literature and will be reviewed in Section VIII.

6. Comparative Evaluation

6.1 Methodology

The eight different credentialing systems are evaluated on 10 dimensions: blockchain platform, storage architecture, token standard, hash algorithm, batch capability, on-chain revocation, access control model, privacy mechanism, W3C VC alignment, and open-source status. The dimensions have been selected based on the priorities of the three major groups of stakeholders: issuers, holders and verifiers.

The comparison data are published system descriptions, technical documentation and, if available, open-source repositories. This is a comparative study of characteristics in a qualitative manner with structure. There was no competing system that was independently deployed, instrumented and benchmarked. Any other performance statistics were self-reported by the authors of the systems.

6.2 Multi-Dimensional Comparison

Table 2: Multi-Dimensional Comparison of Blockchain Credentialing Systems

Dimension	**Blockcerts [16]**	**EduCTX [18]**	**CredChain [19]**	**Alam [20]**	**SuperCert [22]**	**BlockMEDC [23]**	**ZKB AR-V [26]**	**Indy/ Aries [32]**	**Attestify**
Blockchain	Bitcoin	Ark	Ethereum	Ethereum	Ethereum	Eth. L2	Polygon	HL Indy	**Ethereum**
Storage	Off-chain	On-chain	IPFS+chain	IPFS+chain	On-chain	IPFS+L2	On-chain	Off-chain	**IPFS+chain**
Token	None	None	ERC-721	None	None	None	None	VC/VP	**ERC-721 SBT**
Hash	Merkle	N/A	SHA-256	SHA-256	SHA-256	SHA-256	ZK circuit	N/A	**SHA-256+CID**

Batch	✓	✗	✗	✗	✗	✗	✗	✓	✓
Revocation	Hosted	✗	✗	✗	✗	✓	✗	✓	✓ **(on-chain)**
Access Ctrl	Open	Open	Open	Open	Open	Role	Open	DID	**Consortium**
Privacy	✗	✗	✗	Encrypted	✗	✗	**ZKP**	**ZKP**	✗
W3C VC	Partial	✗	✗	✗	✗	✗	✗	✓	Pathway
Open source	✓	✓	✗	✗	✗	✗	✗	✓	✓

6.3 Differentiation Analysis

The surveyed field is divided into five axes when separating of attestify.

SBT-enforced non-transferability. The only system in this comparison that has credentials as non-transferable Soulbound Tokens is the credential system. CredChain [19] is built on top of the standard transferable ERC-721 tokens. All of the other systems surveyed do not provide any tokenisation or have non-token representations. The key value of SBTs is that they prevent credential ownership from being lost once they have been issued, a property that academic credentials should have.

Dual-anchor document integrity. The writing of both a digest SHA-256 sum and an IPFS CID to on-chain storage provides a system with no other surveyed system can match the redundant integrity model. Instead of using a second retrieval and re-verification path, like a content-addressed pointer, CredChain and Alam et al. store lone hash anchors.

Complete lifecycle coverage. All lifecycle stages in Table 1 are available in one framework. Blockcerts is a good approximation but lacks on-chain revocation and identity bound token minting. Such a wide system is not found anywhere else.

Batch issuance. Unlike the other systems in the comparison that are only able to issue a single credential per transaction, Attestify is the only system in the comparison that is Ethereum-based and has developed a practical solution for issuing hundreds or thousands of credentials per academic cycle. While blockcerts aggregates the blocks on Bitcoin, this does not provide on-chain state entries per credential.

Open-source availability. Only Attestify, Blockcerts, EduCTX, and Hyperledger Indy release their full source code, providing for independent audit, extension and institutional deployment.

6.4 Dimensions Where Existing Systems Excel

Honest assessment requires to identify the strengths of the competitors in comparison to Attestify.

Privacy. The selective disclosure of zero knowledge is implemented in ZKBAR-V [26] and Hyperledger Indy [32]. Attestify does not have any verification functionality which preserves privacy. This is its biggest issue, as privacy compliance is becoming a regulatory requirement throughout the European Union and other similar areas.

Gas efficiency. BlockMEDC [23] and ZKBAR-V [26] operate on Layer-2 networks, which operate at about 10x lower gas costs than Attestify on L1. The combined gas usage of Attestify (hash anchoring + SBT minting) on Sepolia is 270,000 to 330,000 gas units, while L2 systems are estimated at 30,000 to 40,000 gas units. Gas Unit counts (in Sepolia) are directly comparable to mainnet as EVM execution is the same, gas pricing isn't.

Standards compliance. Hyperledger Indy is the first implementation of the W3C DID/VC specification. Blockcerts [16] partially aligns with Open Badges v2. Attestify has outlined a clear migration journey (Section VII) but does not yet have W3C VC Data Model conformance.

7. W3C Standards Alignment

W3C Verifiable Credentials Data Model v2.0 [11] is the current leading interoperability standard for the exchange of digital credentials. It structures a verifiable credential into seven essential properties: @context (JSON-LD context declarations), type, issuer (a DID identified entity), credentialSubject, issuanceDate, proof (a Linked Data Proof or equivalent), and credentialStatus.

Attestify is currently offering working alternatives for a subset. The timestamps of issuances are linked to block.timestamp (Unix epoch at block inclusion). The studentId mapping key is used for subject linkage. The on-chain isRevoked flag determines the status of a revocation. The system is missing the use of JSON-LD context declarations, the use of DID based issuer identification, and a conformant proof envelope.

The system would be brought into conformance with four changes, all of which would be in the metadata formatting layer. All credential metadata would be contained within a JSON-LD envelope, that would include the W3C VC context URI. A did:ethr resolver would be able to resolve DID Documents from an Ethereum address, following the W3C DID specification [37]. The existing transaction signature would be converted into an EthereumEip712Signature2021 signature. The revocation status will be brought to the surface via StatusList2021 [35]. None of these changes affects any storage variables, access control logic or token mechanics in the smart contract.

In fact, there is already a credentialing framework for educational institutions in Europe that is aligned with the W3C, namely EBSI [36]. The compatibility with EBSI would lead to integration with the European Digital Identity framework, which would have many implications for adoption in the EU member states.

8. Discussion

8.1 Addressing the Research Questions

The literature survey (Section III) and taxonomy (Section IV) helped to answer **RQ1**. The field has been divided into three generations: hash anachoring on Bitcoin, programmable smart contract platforms on Ethereum, and the latest generation of privacy-preserving and L2 solutions. Fragmentation persists. There is no current solution that provides lifecycle management, SBT identity binding, batch issuance and revocation on-chain. The most significant unresolved constraints are privacy, cost and standards interoperability.

The system description (Section V) and the comparative evaluation (Section VI) answered **RQ2**. For Attestify, there are 5 differentiation axes:

SBT non-transferability, dual-anchor integrity, lifecycle completeness, batch issuance, and open-source availability. These features directly address the gaps in identity-binding and lifecycle that has been documented in the literature. However, in privacy and transaction fees, ZKBAR-V and Hyperledger Indy have definite advantages.

The four open challenges following are the solutions that were used to resolve the challenge **RQ3**. Any progress made on one challenge will probably lead to opportunities and constraints on the rest.

8.2 Open Challenge 1: Privacy Preservation

The biggest drawback of Attestify, and many other public-blockchain-based credentialing systems, is the absence of privacy-preserving verification. Every credential metadata (student ID, institutional name, issuance time) is permanently stored on transparent ledgers and is globally visible. Finally, that exposure contradicts the GDPR principles of data minimisation and the right to erasure (Article 17) [9] and FERPA [8].

The most technically advanced approach to selective disclosure is zero-knowledge constructions. There are also some working implementations, such as ZKBAR-V [26] and Hyperledger Indy [32]. If Attestify supports ZK-SNARK circuits (via Circom or ZoKrates), then credential holders could create brief proofs of credential possession, without revealing the content of the credential, that can be verified on-chain via the SHA-256 anchor. The hybrid storage architecture provides a natural integration surface, the off-chain credential is the private witness, on-chain hash is the public input.

Sedlmeir et al. [28] claimed that privacy-preserving verification is a legal necessity, not an optional added feature, within the EU or similar data protection legislation for any system handling personal educational information.

8.3 Open Challenge 2: Economic Scalability

Whenever it comes to gas fees on Ethereum Layer-1, it's a real obstacle for adoption. The total amount that an institution that issues n credentials, each costing g gas units, spends on gas per year is:

$$C_{L1} = n \cdot g \cdot p$$

With n = 5,000 and g ≈ 300,000, annual gas consumption reaches 1.5×10^9 units. The resulting "Ether expenditure" is significant at current base-fees of the mainnet. It scales linearly. There is zero economies of scale for per transaction issuance and each credential has an equal marginal cost of g · p. There have been several reductions in the availability cost of L2 data, such as blob transactions (EIP-4844) and measured L2 credential deployments [23, 26] that have shown gas savings from 90 to 94 percent.

There are two strategies to be applied together. Moving the smart contract to an L2 rollup (Arbitrum One, Optimism Bedrock, zkSync Era) retains EVM compatibility and involves little or no changes to the smart contract. An alternative route is using Merkle-based batch anchoring. The institution builds a binary Merkle tree on the digests of n credentials off-chain and stores the 32-byte root on-chain. The cost of per-credential amortised storage is now:

$$C_{\text{amort}} = \frac{g_s}{n}$$

In this case, the cost of a single write of a Merkle root is referred to as g_8. As n increases, Camort gets closer to zero and changes linear scaling of cost into almost constant total on-chain costs per batch. To verify the individual credentials, a person builds the Merkle path from the leaf to the published root, outside the blockchain.

8.4 Open Challenge 3: Governance and Key Management

At the present time Attestify has one EOA with full admin rights. That address controls issuer registration

and deregistration unilaterally. This is a single point of compromise, which is not acceptable for multi-institutional production environments. Most of the surveyed literature on Ethereum credentialing systems share this design limitation.

Production-grade governance would require either multi-signature authorisation (a Gnosis Safe multisig with an m-of-n threshold) or a DAO model in which sensitive operations (issuer authorisation, revocation) require on-chain agreement from multiple institutional stakeholders. The administrative key should be stored in a hardware security module (HSM) with key rotation and audit logging. These measures change the security model from protecting a single private key from exfiltration to a Byzantine fault tolerance problem distributed over m independent signers.

8.5 Open Challenge 4: Standards Convergence

Credential standards are still being developed. The W3C VC Data Model [11], Open Badges v3 (now structurally aligned with the VC specification), EBSI [36] and some national digital identity frameworks are evolving along partially overlapping but not yet converged trajectories. Until a critical mass of systems adopt common data models and proof formats, each platform is an isolated trust silo, not a node in a larger verification network.

Section VII showed that it is technically possible to make Attestify conform to W3C standards without architectural changes. But convergence across a whole ecosystem is a problem that no single system can solve on its own. Aligning early with W3C and EBSI specifications is a strategic choice with potential first-mover advantages in institutional adoption.

8.6 Implications for Institutional Adoption

Technically, blockchain credentialing is mature enough for constrained pilots. It's not ready for enterprise scale production. There are four unmet preconditions. L2 migration must be within institutional budgets for each credential. Privacy-preserving selective disclosure has to satisfy GDPR-class regulatory requirements. Multi-party governance must build trust at the consortium level. Standards alignment needs to support portability of credentials across platforms. These are not incremental improvements; they are structural preconditions for adoption at scale.9. Conclusion

This paper reviewed the existing literature on blockchain-based academic credentialing and introduced Attestify as a reference framework that combines capabilities that no existing system has combined. The literature evolves from first-generation Bitcoin-based hash anchoring, to Ethereum smart contract platforms with programmable credential logic, to current privacy-preserving and L2 designs. The hallmark is fragmentation: individual systems tend to address one or two dimensions of credentialing, leaving the rest open.

Attestify fills some of that gap. It bundles together soulbound token non-transferability, dual-anchor document integrity, full lifecycle management, batch issuance, and consortium access control in a single open-source Ethereum deployment. This feature combination was found to be unique in an evaluation against eight systems . It also concluded that there are three dimensions in which ZKBAR-V and Hyperledger Indy still have definitive advantages: privacy, gas efficiency and standards compliance.

Privacy by integrating zero knowledge proofs, economic efficiency by moving to L2 rollups, decentralization of governance by adopting a multi-signature or DAO based structure, and interoperability by conforming to W3CVCandDIdid. Address each one of these issues will require system level technical research and sustained cooperation between the academy, the standardization organizations, the academic institutions and the regulatory bodies.

References

1. A. Grech and A. F. Camilleri, "Blockchain in Education," *JRC Science for Policy Report*, European Commission, 2017. doi: 10.2760/60649.
2. W. Gräther, S. Kolvenbach, R. Ruber, J. Schütte, C. Torres, and F. Wendland, "Blockchain for Education: Lifelong Learning Passport," in *Proc. 1st ERCIM Blockchain Workshop*, pp. 26–32,

2018.
3. P. Bhaskar, C. K. Tiwari, and A. Joshi, "Blockchain in Education Management: Present and Future Applications," *Interactive Technology and Smart Education*, vol. 18, no. 1, pp. 1–17, 2021. doi: 10.1108/ITSE-07-2020-0102.
4. A. M. Antonopoulos and G. Wood, *Mastering Ethereum: Building Smart Contracts and DApps*, 2nd ed., O'Reilly Media, 2022.
5. J. Benet, "IPFS: Content Addressed, Versioned, P2P File System," *arXiv:1407.3561*, 2014.
6. E. G. Weyl, P. Ohlhaver, and V. Buterin, "Decentralized Society: Finding Web3's Soul," *SSRN Electronic Journal*, May 2022. doi: 10.2139/ssrn.4105763.
7. G. Wood, "Ethereum: A Secure Decentralised Generalised Transaction Ledger," *Ethereum Yellow Paper*, 2024 revision.
8. U.S. Department of Education, "Family Educational Rights and Privacy Act (FERPA)," *20 U.S.C. § 1232g*, 1974.
9. European Parliament, "General Data Protection Regulation (GDPR)," *Regulation (EU) 2016/679*, 2016.
10. A. Voskobojnikov, O. Wiese, M. Kober, D. Ermakova, and G. Balapour, "Surviving the Cryptojungle: Perception and Management of Risk Among North American Cryptocurrency Users," *Financial Innovation*, vol. 7, no. 1, pp. 1–24, 2021.
11. W3C, "Verifiable Credentials Data Model v2.0," *W3C Recommendation*, Mar. 2024. [Online]. Available: https://www.w3.org/TR/vc-data-model-2.0/
12. M. Alonso, M. Arenas, and P. Fernandez, "Blockchain-Based Credential Verification: A Systematic Review," *Education and Information Technologies*, vol. 29, pp. 4935–4974, 2024. doi: 10.1007/s10639-023-12032-1.
13. S. Morshed, B. Siddiqui, and M. S. Rahman, "Blockchain in Education: A Comprehensive Survey on Applications, Challenges, and Opportunities," *IEEE Access*, vol. 12, pp. 22612–22636, 2024.
14. A. R. Hevner, S. T. March, J. Park, and S. Ram, "Design Science in Information Systems Research," *MIS Quarterly*, vol. 28, no. 1, pp. 75–105, 2004.
15. K. Peffers, T. Tuunanen, M. A. Rothenberger, and S. Chatterjee, "A Design Science Research Methodology for Information Systems Research," *Journal of Management Information Systems*, vol. 24, no. 3, pp. 45–77, 2007.
16. J. P. Schmidt, "Digital Certificates Project," *MIT Media Lab*, 2017. [Online]. Available: https://www.blockcerts.org
17. M. Sharples and J. Domingue, "The Blockchain and Kudos: A Distributed System for Educational Record, Reputation and Reward," in *Proc. European Conf. Technology Enhanced Learning (EC-TEL)*, pp. 490–496, 2016.
18. M. Turkanović, M. Hölbl, K. Košić, M. Heričko, and A. Kamišalić, "EduCTX: A Blockchain-Based Higher Education Credit Platform," *IEEE Access*, vol. 6, pp. 5112–5127, 2018. doi: 10.1109/ACCESS.2018.2789929.
19. M. Ammad-Ud-Din *et al.*, "CredChain: Academic and Professional Certificate Verification System using Blockchain," in *Proc. IEEE Intl. Conf. Blockchain*, 2024.
20. T. Alam, A. Qayyum, and M. Benaida, "IPFS-Blockchain Smart Contracts Based Conceptual Framework to Reduce Certificate Frauds in the Academic Field," *Information*, vol. 14, no. 8, p. 446, 2023. doi: 10.3390/info14080446.
21. R. Arenas and P. Fernandez, "CredTrust: A Multidimensional Framework for Evaluating Trust in Blockchain-Based Educational Credentials," in *Proc. IEEE Frontiers in Education (FIE)*, 2023.
22. A. R. Khan, Q. Rabbani, and A. Iqbal, "SuperCert: An Anti-Fraud Identity Intelligence Blockchain Solution for Educational Certificates," *IEEE Access*, vol. 12, pp. 34208–34222, 2024. doi: 10.1109/ACCESS.2024.3412156.
23. M. El-Hajj, A. Fadlallah, M. Chamoun, and A. Serhrouchni, "BlockMEDC: Blockchain Smart Contracts for Securing Moroccan Higher Education Digital Certificates," *IEEE Access*, vol. 12, pp. 41983–41996, 2024. doi: 10.1109/ACCESS.2024.3378095.
24. Ethereum Foundation, "EIP-4844: Shard Blob Transactions (Proto-Danksharding)," *Ethereum Improvement Proposals*, Mar. 2024.
25. Simply Staking, "Layer-2 Ecosystem Report: Post-Dencun Performance Analysis," 2024.

[Online]. Available: https://www.simplystaking.com
26. ZKBAR-V Authors, "Zero-Knowledge Proof-Enabled Blockchain-Based Academic Record Verification System," *ResearchGate Preprint*, 2024.
27. A. Mühle, A. Grüner, T. Gayvoronskaya, and C. Meinel, "A Survey on Essential Components of a Self-Sovereign Identity," *Computer Science Review*, vol. 30, pp. 80–86, 2018. doi: 10.1016/j.cosrev.2018.10.002.
28. J. Sedlmeir, R. Smethurst, A. Rieger, and G. Fridgen, "Digital Identities and Verifiable Credentials," *Business and Information Systems Engineering*, vol. 63, pp. 603–613, 2021. doi: 10.1007/s12599-021-00722-y.
29. T. J. Wilson, "ERC-5192: Minimal Soulbound NFTs," *Ethereum Improvement Proposals*, Jul. 2022. [Online]. Available: https://eips.ethereum.org/EIPS/eip-5192
30. S. Park, B. Kim, and H. Kim, "Soulbound Tokens: Enabler for Privacy-Aware and Decentralized Authentication Mechanism in Medical Data Storage," *Sensors*, vol. 23, no. 18, p. 7930, 2023. doi: 10.3390/s23187930.
31. C. Li, B. Palanisamy, and R. Liu, "SBT-Based Access Control: Implementing Non-Transferable Tokens for Decentralized Identity," in *Proc. IEEE Intl. Conf. Blockchain and Cryptocurrency*, 2023.
32. Hyperledger Foundation, "Hyperledger Indy and Aries: Self-Sovereign Identity Infrastructure," 2023. [Online]. Available: https://www.hyperledger.org/projects/hyperledger-indy
33. A. Preukschat and D. Reed, *Self-Sovereign Identity: Decentralized Digital Identity and Verifiable Credentials*, Manning Publications, 2021.
34. uPort, "ERC-1056: Ethereum Lightweight Identity," *Ethereum Improvement Proposals*, 2018.
35. W3C, "Status List 2021," *W3C Community Group Report*, 2023. [Online]. Available: https://www.w3.org/TR/vc-status-list/
36. European Commission, "European Blockchain Services Infrastructure (EBSI)," 2024. [Online]. Available: https://ec.europa.eu/digital-building-blocks/wikis/display/EBSI
37. W3C, "Decentralized Identifiers (DIDs) v1.0," *W3C Recommendation*, Jul. 2022. [Online]. Available: https://www.w3.org/TR/did-core/
38. OpenZeppelin, "Contracts v5.0: ERC721, Ownable, ReentrancyGuard," 2024. [Online]. Available: https://docs.openzeppelin.com/contracts/5.x/