International Journal of Creative and Open Research in Engineering and Management
ISSN: **3108-1754 (Online)**
Volume 02 Issue 08 August-2026 | Impact Factor: 3.5

Adaptive ZKP Selection Framework for Privacy-Preserving Decentralized Identity Using DID and Verifiable Credentials

Sanchita Shukla
Department of Computer Science & Engineering (AI & ML)
KIPM College of Engineering & Technology, GIDA, Gorakhpur

How to Cite this Article:

Shukla, S. (2026). Adaptive ZKP Selection Framework for Privacy-Preserving Decentralized Identity Using DID and Verifiable Credentials. International Journal of Creative and Open Research in Engineering and Management, 02(8), 1-9. https://doi.org/10.55041/ijcope.v2i8.249

https://doi.org/10.55041/ijcope.v2i8.249

Abstract

The rapid growth of digital services has increased the demand for identity systems that provide strong authentication while minimizing unnecessary disclosure of personal information. Conventional identity management architectures generally depend on centralized identity providers and frequently require users to disclose complete identity attributes even when a service requires only a limited assertion. Decentralized Identifiers (DIDs) and Verifiable Credentials (VCs) provide an alternative model in which identity holders can manage cryptographically verifiable credentials independently of a centralized identity provider. However, the privacy guarantees of decentralized identity systems depend substantially on the cryptographic proof mechanism used during credential presentation.

Zero-Knowledge Proofs (ZKPs) enable a prover to demonstrate knowledge of a secret or the validity of a statement without revealing the underlying secret. Different ZKP and selective-disclosure mechanisms exhibit significantly different characteristics with respect to proof size, generation time, verification time, communication overhead, computational requirements, privacy guarantees, interoperability, and implementation complexity. Consequently, selecting a single proof mechanism for every decentralized identity scenario can result in unnecessary computational cost or inadequate privacy protection.

This paper proposes an **Adaptive ZKP Selection Framework (AZSF)** for privacy-preserving decentralized identity using DIDs and Verifiable Credentials. The framework dynamically selects an appropriate proof mechanism according to the privacy sensitivity of requested attributes, disclosure requirements, verifier trust conditions, computational resources, proof-generation latency, communication constraints, interoperability requirements, and unlinkability requirements. The proposed architecture introduces a policy-driven decision layer between the credential wallet and proof-generation subsystem. It evaluates candidate mechanisms including selective-disclosure signatures, BBS-based proofs, SD-JWT-based selective disclosure, and general-purpose succinct zero-knowledge proof systems such as zk-SNARK/PLONK-style approaches.

A formal multi-criteria decision model is developed to represent the selection process. The framework defines privacy, performance, communication, interoperability, and deployment criteria and computes an adaptive suitability score for each candidate proof mechanism. A threat model covering credential theft, replay, correlation, malicious verifiers, issuer compromise, metadata leakage, and proof substitution is presented. The paper further proposes an experimental methodology for evaluating proof generation time, verification time, proof size, communication overhead, privacy leakage, unlinkability, and resource consumption. The proposed framework provides a systematic foundation for choosing cryptographic proof mechanisms according to application requirements instead of adopting a one-size-fits-all approach.

Index Terms— Decentralized Identity, Decentralized Identifiers, Verifiable Credentials, Zero-Knowledge Proofs, Privacy-Preserving Identity, Selective Disclosure, BBS Signatures, SD-JWT, Self-Sovereign Identity, Privacy Engineering, Adaptive Cryptography.

I. INTRODUCTION

Digital identity has become a fundamental component of modern computing systems. Online education, financial services, healthcare, e-governance, employment, travel, and enterprise applications increasingly require users to establish claims about their identity, qualifications, authorization, age, residence, or organizational affiliation. Conventional identity systems generally depend on centralized identity providers, federated authentication systems, or service-specific databases. Although these models can provide efficient authentication, they may create significant privacy risks because identity providers and relying parties can accumulate extensive records of users' activities.

A fundamental privacy problem is **over-disclosure**. For example, a service may require proof that a user is above a specified age, while a conventional identity document may expose the user's complete date of birth, name, address, identification number, photograph, and other information. Similarly, an educational service may require proof that an individual holds a particular qualification but may receive substantially more personal information than is necessary for that purpose.

Decentralized Identity (DI) attempts to address this problem by allowing identity subjects to use decentralized identifiers and cryptographically verifiable credentials. The World Wide Web Consortium (W3C) Decentralized Identifiers specification defines DIDs as identifiers designed to be decoupled from centralized registries, identity providers, and certificate authorities. A DID can resolve to a DID Document containing verification methods and service information that support cryptographic control and interaction with the identifier [1]. (W3C)

Verifiable Credentials extend this architecture by allowing issuers to make cryptographically verifiable claims about subjects. The W3C Verifiable Credentials Data Model 2.0 defines an issuer-holder-verifier ecosystem and explicitly recognizes selective disclosure and unlinkable disclosure as privacy-related concepts [2]. (W3C)

Zero-Knowledge Proofs (ZKPs) provide an additional privacy layer. NIST describes ZKPs as a major privacy-enhancing cryptographic mechanism that permits a party to establish the truth of a mathematical statement without revealing unnecessary information associated with the proof [3]. (NIST Computer Security Resource Center)

Despite these advances, a significant engineering challenge remains: **which proof mechanism should be selected for a particular identity-verification request?**

A university admission system may prioritize interoperability and low verification latency. A financial application may require strong unlinkability and predicate proofs. A resource-constrained mobile device may prioritize low proof-generation cost. A high-assurance government service may prioritize cryptographic expressiveness even if the proof-generation process is computationally expensive.

Therefore, the central argument of this paper is that:

The cryptographic proof mechanism in a decentralized identity system should be selected adaptively according to the privacy, performance, interoperability, and resource requirements of the verification transaction.

This paper proposes an **Adaptive ZKP Selection Framework (AZSF)** to address this problem.

II. RESEARCH MOTIVATION AND PROBLEM STATEMENT

Existing decentralized identity architectures typically separate identity management from cryptographic proof generation. Although this modularity is beneficial, implementations frequently select a proof technology at the architecture-design stage rather than dynamically according to application requirements.

This creates five important problems.

A. One-Size-Fits-All Cryptography

Different applications have different requirements. A mechanism optimized for compact selective disclosure may not be optimal for complex predicates involving multiple attributes.

B. Privacy-Performance Trade-Off

Increasing cryptographic privacy can increase computational cost, proof-generation latency, memory requirements, or proof size.

C. Heterogeneous Devices

Identity wallets can operate on desktop computers, smartphones, embedded devices, browser environments, or hardware-secured systems. These platforms have different computational capabilities.

D. Interoperability

DID and VC ecosystems contain multiple credential formats and proof suites. W3C's current VC ecosystem includes multiple approaches to securing and selectively disclosing credential information. (W3C)

E. Privacy Is Multi-Dimensional

Privacy cannot be represented simply as "data is encrypted." Important properties include:

- data minimization,
- selective disclosure,
- unlinkability,
- holder binding,
- resistance to correlation,
- resistance to replay,
- issuer privacy,
- verifier privacy,
- metadata minimization.

Therefore, a mechanism-selection framework must evaluate privacy as a multidimensional property.

III. RESEARCH OBJECTIVES

The primary objective of this research is to develop an adaptive decision framework for selecting an appropriate zero-knowledge or selective-disclosure proof mechanism for decentralized identity applications.

The specific objectives are:

1. To design a privacy-preserving DID and VC architecture incorporating adaptive proof selection.
2. To identify the cryptographic, privacy, performance, and interoperability characteristics relevant to proof selection.
3. To develop a formal multi-criteria scoring model for selecting proof mechanisms.
4. To support selective disclosure and predicate-based credential verification.
5. To minimize unnecessary disclosure of personal information.
6. To incorporate device and network constraints into cryptographic selection.
7. To evaluate candidate mechanisms using measurable performance and privacy metrics.
8. To provide a reproducible experimental methodology for future implementation.
9. To analyze security threats against the proposed architecture.

10. To establish guidelines for selecting cryptographic proof mechanisms according to application context.

IV. RESEARCH QUESTIONS

The proposed research investigates the following questions.

RQ1: How can a decentralized identity system dynamically select a cryptographic proof mechanism according to privacy and application requirements?

RQ2: Which parameters have the greatest influence on proof-mechanism selection?

RQ3: How can privacy requirements such as selective disclosure and unlinkability be incorporated into a formal decision model?

RQ4: What are the performance trade-offs among BBS-based proofs, SD-JWT-based selective disclosure, and general-purpose ZKP systems?

RQ5: Can adaptive proof selection reduce unnecessary computational and communication overhead while preserving required privacy guarantees?

RQ6: How can the framework maintain interoperability with emerging DID and VC standards?

V. CONTRIBUTIONS OF THE PAPER

The principal contributions are as follows:

1. Adaptive Proof Selection Model

A formal decision model is proposed for selecting a proof mechanism based on application-specific requirements.

2. Privacy-Aware Decision Layer

The framework explicitly considers privacy properties such as selective disclosure and unlinkability instead of treating cryptographic performance as the only optimization criterion.

3. Multi-Criteria Evaluation Function

A weighted suitability function is developed using privacy, performance, communication, interoperability, and resource criteria.

4. DID–VC–ZKP Architecture

A complete conceptual architecture is proposed connecting DID resolution, credential management, policy evaluation, adaptive proof selection, proof generation, and verification.

5. Security and Threat Model

The paper defines threats affecting decentralized identity and evaluates how the proposed architecture addresses them.

6. Reproducible Experimental Methodology

A benchmark methodology is proposed to compare candidate proof mechanisms without claiming experimental results that have not yet been generated.

VI. BACKGROUND

A. Decentralized Identifiers

A DID is a decentralized identifier associated with a subject and designed to enable cryptographic control without requiring a centralized identity provider. DID Documents can contain verification methods, authentication relationships, and service endpoints [1]. (W3C)

A simplified DID structure can be represented as:

$$DID = did:method:methodSpecificId$$

The DID resolution process maps a DID to a DID Document:

$$Resolve(DID) \rightarrow DIDDocument$$

The DID Document can provide the public verification material required to authenticate cryptographic operations.

The W3C DID specification intentionally does not mandate one underlying technological infrastructure. DID methods may use distributed ledgers, decentralized storage, peer-to-peer systems, or other architectures [1]. (W3C)

B. Verifiable Credentials

A Verifiable Credential is a tamper-evident digital credential containing one or more claims made by an issuer.

The basic ecosystem contains three actors:

1. **Issuer**
2. **Holder**
3. **Verifier**

The issuer generates a credential:

$$VC = Sign_{SK_I}(Claims)$$

where $SK_I$is the issuer's private key.

The holder stores the credential and later generates a presentation satisfying a verifier request.

A simplified presentation model is:

$$VP = Present(V, C, PolicyProof)$$

The verifier evaluates:

$$Verify(V, P, PK_IPolicy) \rightarrow \{T, rueFalse\}$$

The W3C VC Data Model 2.0 explicitly supports verifiable presentations and identifies selective and unlinkable disclosure as privacy concepts [2]. (W3C)

VII. ZERO-KNOWLEDGE PROOFS

A Zero-Knowledge Proof allows a prover $P$to convince a verifier Vthat a statement $x$is true without revealing the secret witness w.

A proof system can be represented as:

$$Setup \rightarrow (p, kvk)$$
$$Proof = Prove(p, k, xw)$$
$$Verify(v, k, xProof) \rightarrow \{0, 1\}$$

A secure ZKP generally seeks three properties:

Completeness

An honest prover should convince an honest verifier.

$$Pr\big[Verify\big(x, Prove(x, w)\big) = 1\big] \approx 1$$

Soundness

A dishonest prover should not be able to prove a false statement except with negligible probability.

Zero-Knowledge

The verifier should learn no information about the witness beyond what is implied by the statement.

NIST identifies ZKPs as a major component of privacy-enhancing cryptography [3]. (NIST Computer Security Resource Center)

VIII. SELECTIVE DISCLOSURE IN VERIFIABLE CREDENTIALS

Selective disclosure allows the holder to disclose only the claims necessary for a particular transaction.

Assume a credential contains:

$$C = \{N, a, m, e, D, OBAddressNationalityDegreeUniversity\}$$

A verifier may only require:

$$R = \{Degree\}$$

A privacy-preserving presentation should disclose:

$$Present(C, R) = \{D, egreeProof\}$$

rather than:

$$Present(C) = \{N, a, m, e, D, OBAddressNationalityDegreeUniversity\}$$

This implements the privacy principle of data minimization.

The current W3C VC ecosystem includes cryptosuites supporting selective disclosure, including ECDSA-based selective disclosure and BBS-based mechanisms. (W3C)

IX. EXISTING CRYPTOGRAPHIC APPROACHES

The proposed framework does not assume that one mechanism is universally superior. Instead, it treats each mechanism as a candidate with different characteristics.

A. SD-JWT

Selective Disclosure JWT (SD-JWT) provides selective disclosure for JSON Web Tokens using salted hashes. Instead of placing selectively disclosable claims directly in the issuer-signed payload, digests can represent the claims, while the holder later provides selected disclosures [4]. (RFC Editor)

Advantages include:

- strong compatibility with JWT infrastructure;
- comparatively straightforward deployment;
- support for selective disclosure;
- suitability for web-oriented ecosystems.

Potential limitations include:

- privacy properties depend on presentation construction;
- selective disclosure is different from full general-purpose ZKP expressiveness;
- correlation considerations must be carefully evaluated.

B. BBS-Based Proofs

BBS-based mechanisms are particularly relevant to privacy-preserving Verifiable Credentials.

The W3C BBS Cryptosuite provides selective disclosure and unlinkable derived proofs [5]. (W3C)

Conceptually:

$$Signature = BBS.Sign(S, K_I M)$$

The holder can derive a proof for a subset:

$$P = BBS.ProofGen(S, ignatureS)$$

where $S \subseteq M$.

The verifier can verify:

$$BBS.ProofVerify(P, P, K_I S) = True$$

An important property is that the holder can produce different derived proofs from the same credential. The W3C BBS specification describes these derived proofs as unlinkable to the original signature under the intended cryptographic model [5]. (W3C)

However, BBS remains a specialized cryptographic mechanism, and its implementation and interoperability requirements must be considered carefully. Recent standardization work also demonstrates that the security analysis of BBS-based systems remains an active research area. (arXiv)

C. General-Purpose zk-SNARK/PLONK-Style Systems

General-purpose succinct zero-knowledge systems can express more complex statements than simple attribute disclosure.

For example, a verifier may require:

$$Age \geq 18$$

without learning the exact age.

A more complex predicate might be:

$$Age \geq 18 \wedge Country = IN \wedge CredentialValid = True$$

Such predicates can be represented as arithmetic circuits:

$$C(w, x) = 1$$

and proven using a suitable ZKP protocol.

The major advantages are expressive predicates and strong privacy capabilities. The disadvantages may include greater proof-generation complexity, setup or circuit-management requirements depending on the system, and higher computational cost.

X. RESEARCH GAP

Existing research has investigated:

- decentralized identity;
- Self-Sovereign Identity;
- Verifiable Credentials;
- selective disclosure;

- zero-knowledge authentication;
- blockchain-based identity;
- ZKP-based credential verification;
- BBS signatures;
- SD-JWT.

However, the literature and standards landscape does not eliminate the engineering problem of **runtime proof-mechanism selection based on application context**.

A recent study of cryptographic mechanisms for selective disclosure highlights that multiple cryptographic approaches exist, including commitment-based mechanisms and non-interactive ZKP approaches [6]. (ScienceDirect)

Similarly, threat-modeling research emphasizes that decentralized identity systems should be evaluated against multiple security and trust threats rather than relying solely on architectural decentralization [7]. (Springer)

Therefore, this paper addresses a specific gap:

There is a need for a policy-driven framework that evaluates the requirements of a credential-verification transaction and dynamically selects an appropriate privacy-preserving proof mechanism.

XI. PROPOSED ADAPTIVE ZKP SELECTION FRAMEWORK

A. Architectural Overview

The proposed architecture consists of seven logical layers:

1. DID Layer
2. Credential Layer
3. Policy Layer
4. Adaptive Selection Layer
5. Proof Generation Layer
6. Verification Layer
7. Audit and Privacy Monitoring Layer

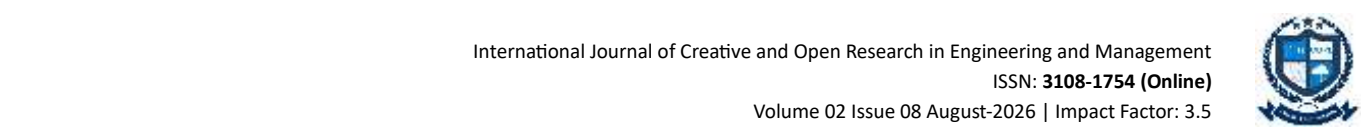

The conceptual workflow is:

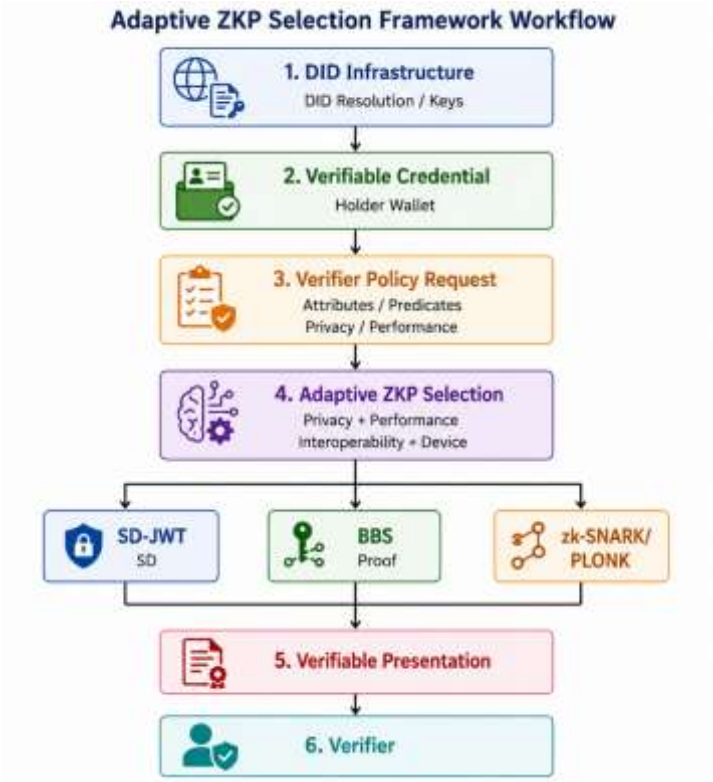

XII. ADAPTIVE SELECTION MODEL

Let the candidate proof mechanisms be:

$$M = \{m_1, m_2, \dots, m_n\}$$

For example:

$$M = \{S, D, JWTBBSZKSNARK\}$$

Let the selection criteria be:

$$C = \{P, r, i, v, a, cyPerformanceSizeInteroperabilityResourceExpressiveness\}$$

Each mechanism receives a normalized score between 0 and 1.

The suitability score for mechanism $m_i$is:

$$Score(m_i) = \sum_{j=1}^{k} w_j \, S_{ij}$$

where:

$$\sum_{j=1}^{k} w_j = 1$$

and:

- w_j= importance weight of criterion j;

- S_{ij}= normalized score of mechanism $i$for criterion j.

The selected mechanism is:

$$m^* = \arg\max_{m_i \in M} Score(m_i)$$

However, a pure weighted average is insufficient because some requirements are mandatory.

Therefore, the proposed framework introduces **hard constraints**.

For example:

$$Privacy(m_i) \geq P_{min}$$
$$Latency(m_i) \leq L_{max}$$
$$ProofSize(m_i) \leq S_{max}$$

Only mechanisms satisfying all mandatory constraints enter the weighted ranking.

Thus:

$$M' = \{m_i \in M \mid Constraints(m_i) = True\}$$

and:

$$m^* = \arg\max_{m_i \in M'} Score(m_i)$$

This prevents a mechanism with excellent performance but inadequate privacy from being selected.

XIII. SELECTION CRITERIA

A. Privacy Score

Privacy score can incorporate:

- selective disclosure;
- predicate proving;
- unlinkability;
- holder binding;
- correlation resistance;
- metadata minimization.

Define:

$$P = \alpha_1 SD + \alpha_2 UL + \alpha_3 PR + \alpha_4 HB$$

where:

- SD= selective disclosure;
- UL= unlinkability;
- PR= predicate-proof capability;
- HB= holder binding.

B. Performance Score

Performance can be measured through:

$$Perf = \beta_1(1 - T_g) + \beta_2(1 - T_v) + \beta_3(1 - CPU) + \beta_4(1 - MEM)$$

where:

- T_g= proof-generation time;
- T_v= proof-verification time;
- CPU = normalized CPU utilization;
- MEM = normalized memory consumption.

Values are normalized before aggregation.

C. Communication Score

Let:

$$Comm = 1 - \frac{ProofSize}{MaxProofSize}$$

A smaller proof results in a higher communication score.

D. Interoperability Score

Interoperability evaluates:

- standards support;
- VC compatibility;
- DID ecosystem compatibility;
- wallet compatibility;
- verifier implementation availability;
- deployment maturity.

E. Expressiveness Score

Expressiveness measures whether the mechanism can prove:

- equality;
- inequality;
- range predicates;
- set membership;
- credential validity;
- compound predicates.

For example:

$$Age \geq 18$$

is more expressive than merely revealing an age attribute.

XIV. ADAPTIVE POLICY ENGINE

The policy engine receives a verifier request:

$$R = \{A, t, t, r, i, butesPredicatesPrivacyLatencySizeTrust\}$$

Example:

Verifier:

Required claim: Age >= 18

Required issuer: Government

Maximum latency: 500 ms

Maximum proof size: 50 KB

Unlinkability: Required

Selective disclosure: Required

The policy engine transforms the request into normalized parameters:

$$R' = (P, L, S, U, I, E)$$

The adaptive selector evaluates each candidate mechanism.

Algorithm 1: Adaptive Proof Selection

Input:

Candidate mechanisms M

Verification request R

Device profile D

Network profile N

Output:

Selected mechanism m*

1. Parse verifier policy R
2. Extract privacy requirements
3. Extract performance requirements
4. Extract communication constraints
5. Extract interoperability requirements
6. Obtain device capabilities D
7. Obtain network conditions N
8. Remove mechanisms violating hard constraints
9. Normalize remaining criterion scores
10. Calculate weighted suitability score
11. Apply privacy threshold

12. Select mechanism with maximum score

13. Generate proof using selected mechanism

14. Attach verifier challenge/domain binding

15. Return Verifiable Presentation

XV. CONTEXT-AWARE ADAPTATION

A major feature of AZSF is context awareness.

The same credential may require different proof mechanisms in different environments.

Scenario 1: University Admission

Requirements:

- degree verification;
- issuer authenticity;
- low computational overhead;
- high interoperability.

A selective-disclosure mechanism may be preferred.

Scenario 2: Age Verification

Requirement:

$$Age \geq 18$$

The verifier does not need the exact date of birth.

A predicate-capable ZKP can be preferred.

Scenario 3: Financial Eligibility

Requirements:

$$Income > Threshold$$
$$Residency = India$$
$$CredentialValid = True$$

A more expressive ZKP may be preferred.

Scenario 4: Low-Power Device

If the holder uses a low-resource device:

$$CPU_{available} \ll CPU_{desktop}$$

the selector should penalize computationally expensive mechanisms.

Scenario 5: High-Privacy Environment

If unlinkability is mandatory:

$$UL_{required} = 1$$

mechanisms that cannot satisfy unlinkability are eliminated rather than simply receiving a lower score

XVI. PRIVACY ENGINEERING MODEL

The proposed system incorporates privacy engineering principles throughout the identity lifecycle.

A. Data Minimization

Only required attributes are disclosed.

$$D_{shared} \subseteq D_{available}$$

The objective is:

$$min(|\ D_{shared}\ |)$$

subject to verification correctness.

B. Purpose Limitation

The verifier's request should specify the intended purpose.

For example:

Purpose = "Age Verification"

A credential presentation generated for one purpose should not automatically be reusable for another.

C. Unlinkability

Suppose a holder interacts with verifiers $V_1, V_2, \ldots, V_n$.

The objective is to reduce the probability that two presentations can be correlated:

$$Pr\left[Link\left(V, P_i V P_j\right)\right] \rightarrow 0$$

subject to the system's intended accountability requirements.

BBS-based derived proofs are specifically designed to support selective disclosure and unlinkable presentations [5]. (W3C)

XVII. SECURITY THREAT MODEL

The proposed system assumes the following adversaries.

A. Malicious Verifier

A verifier attempts to collect unnecessary identity attributes.

Countermeasure: selective disclosure and policy enforcement.

B. External Network Attacker

An attacker observes communication.

Countermeasure: encrypted transport, challenge-response, and minimum disclosure.

C. Replay Attacker

An attacker reuses an earlier presentation.

Countermeasure:

$$Challenge = RandomNonce$$

and domain/audience binding.

D. Credential Theft

An attacker obtains a credential from the holder's device.

Countermeasures:

- wallet encryption;
- hardware-backed keys;
- biometric/device authentication;
- holder binding.

E. Issuer Key Compromise

An attacker compromises issuer signing keys.

Countermeasures:

- key rotation;
- DID document updates;
- revocation/status mechanisms;
- trust-list governance.

F. Correlation Attack

Multiple verifiers attempt to determine whether presentations originate from the same credential holder.

Countermeasure: unlinkable presentation mechanisms where required.

G. Metadata Leakage

Even when credential contents remain hidden, timing, identifiers, endpoints, or repeated values can enable correlation.

Countermeasure: privacy-aware protocol design and metadata minimization

XVIII. SECURITY REQUIREMENTS

The proposed framework defines the following requirements.

SR1: Credential Integrity

An attacker must not be able to modify credential claims without detection.

SR2: Proof Soundness

An invalid claim must not be accepted except with negligible probability.

SR3: Zero Knowledge

The proof should reveal no unnecessary information about hidden witness values.

SR4: Selective Disclosure

The holder must be able to disclose only required attributes when supported.

SR5: Replay Resistance

A presentation must be bound to a verifier-generated challenge.

SR6: Holder Binding

Where required, the presentation should demonstrate control of the authorized holder key.

SR7: Unlinkability

When unlinkability is required, repeated presentations should not contain stable correlators unnecessarily.

SR8: Revocation Awareness

The verifier must be able to determine whether the issuer still stands behind a credential when revocation/status checking is required.

XIX. COMPARATIVE FRAMEWORK

The following table represents the **conceptual characteristics that the proposed experimental evaluation should measure** rather than fabricated benchmark results.

Parameter	**Requirement**	**Parameter**	**Requirement**
Selective disclosure	Mandatory	Selective disclosure	Mandatory
Unlinkability	Mandatory	Unlinkability	Mandatory
Predicate proof	Preferred	Predicate proof	Preferred
Proof size	< 50 KB	Proof size	< 50 KB
Verification latency	< 300 ms	Verification latency	< 300 ms
Device	Smartphone	Device	Smartphone
Parameter	**Requirement**	**Parameter**	**Requirement**
Selective disclosure	Mandatory	Selective disclosure	Mandatory
Unlinkability	Mandatory	Unlinkability	Mandatory
Predicate proof	Preferred	Predicate proof	Preferred
Proof size	< 50 KB	Proof size	< 50 KB
Verification latency	< 300 ms	Verification latency	< 300 ms
Device	Smartphone	Device	Smartphone

These values are **qualitative design expectations** and must be replaced or supplemented by measured results in an empirical version of the study.

XX. EXPERIMENTAL METHODOLOGY

To validate AZSF, a prototype should be implemented using at least three candidate mechanisms.

Test Environment

A reproducible evaluation should report:

- CPU model;
- RAM;
- operating system;
- programming language;
- cryptographic library;
- compiler/runtime version;
- device class;
- network characteristics.

Test Variables

Credential sizes:

$$C = \{5, 10, 25, 50, 100\}$$

attributes.

Disclosure ratios:

$$D = \{10, \%, 25, \%, 50\%75\%100\%\}$$

Concurrent verification loads:

$$L = \{1, 10, 50, 100, 500\}$$

The evaluation should measure:

1. proof-generation time;
2. proof-verification time;
3. proof size;
4. credential size;
5. CPU utilization;
6. memory consumption;
7. communication overhead;
8. energy consumption where measurable;
9. privacy leakage;
10. unlinkability;
11. failure rate;
12. interoperability.

XXI. PERFORMANCE METRICS

A. Proof Generation Time

$$T_g = t_{finish} - t_{start}$$

Measured across repeated trials:

$$\bar{T}_g = \frac{1}{n}\sum_{i=1}^{n} T_{g,i}$$

B. Verification Time

$$T_v = t_{verification,end} - t_{verification,start}$$

C. Proof Size

$$S_p = | \, Proof \, |$$

in bytes.

D. Communication Overhead

$$O_c = \frac{S_{presentation} - S_{minimum}}{S_{minimum}}$$

E. Privacy Leakage

Define a privacy leakage index:

$$L_p = \sum_i \lambda_i D_i$$

where $D_i$represents an observable data element and $\lambda_i$represents its privacy sensitivity.

The optimization objective is:

$$min(L_p)$$

subject to:

$$Verify(Presentation) = True$$

XXII. ADAPTIVE SELECTION EXAMPLE

Assume the following verifier policy:

Parameter	**Requirement**
Selective disclosure	Mandatory
Unlinkability	Mandatory
Predicate proof	Preferred
Proof size	< 50 KB
Verification latency	< 300 ms

The framework first eliminates mechanisms violating mandatory requirements.

The remaining mechanisms are scored:

$$Score(m) = 0.30P + 0.20Perf + 0.15Comm + 0.15Inter + 0.10Expr + 0.10Resource$$

Suppose a different verifier has no unlinkability requirement but requires strong JWT interoperability. The weights change:

$$Score(m) = 0.20P + 0.20Perf + 0.25Inter + 0.15Comm + 0.10Expr + 0.10Resource$$

Consequently, the selected mechanism can change according to context.

This is the central principle of AZSF:

The framework optimizes for the verification policy, rather than assuming that one cryptographic mechanism is optimal for every application.

XXIII. PROPOSED SYSTEM WORKFLOW

The complete workflow is:

Step 1: DID Resolution

The verifier resolves issuer and holder DIDs.

Step 2: Credential Discovery

The holder's wallet identifies credentials capable of satisfying the verifier request.

Step 3: Policy Extraction

The verifier's request is converted into privacy and performance constraints.

Step 4: Context Acquisition

The wallet determines:

device capability; network condition; available cryptographic implementations; credential type; requested predicates.

Step 5: Candidate Generation

The selector generates candidate mechanisms.

Step 6: Constraint Filtering

Mechanisms violating mandatory requirements are removed.

Step 7: Multi-Criteria Ranking

Remaining mechanisms are scored.

Step 8: Proof Generation

The selected mechanism produces a proof or selectively disclosed presentation.

Step 9: Challenge Binding

The presentation is bound to:

verifier; transaction; nonce; domain; expiration.

Step 10: Verification

The verifier validates the presentation.

Step 11: Privacy Audit

The system records only privacy-safe operational metadata required for auditing.

XXIV. FORMAL OPTIMIZATION PROBLEM

The selection problem can be represented as:

$$\max_{m \in M} [w_P P(m) + w_F F(m) + w_C C(m) + w_I I(m) + w_E E(m)]$$

subject to:

$$P(m) \geq P_{min}$$
$$T_v(m) \leq T_{max}$$
$$S_p(m) \leq S_{max}$$
$$R(m) \leq R_{max}$$

where:

- P= privacy;
- F= performance;
- C= communication efficiency;

- I= interoperability;
- E= expressiveness;
- R= resource requirement.

This formulation converts proof selection from an informal engineering choice into a measurable optimization problem

XXV. MACHINE-LEARNING EXTENSION

Because the research area includes Artificial Intelligence and Machine Learning, AZSF can be extended with a learning-based selector.

Let:

$$X = [P, L, S, D, N, C, T]$$

where:

- P= privacy requirement; L= latency requirement; S= proof-size constraint; D= device capability; N= network state; C= credential complexity; T= trust environment.

The model learns:

$$f(X) \rightarrow m^*$$

Possible models include:

- Decision Tree; Random Forest; Gradient Boosting; Multi-Layer Perceptron; Contextual Bandit.

However, a learning-based model should not override hard security constraints.

Therefore:

$$MLSelection \subseteq SecurityFeasibleSet$$

This ensures that machine learning optimizes among acceptable cryptographic choices rather than making unconstrained security decisions.

XXVI. PRIVACY-AWARE REWARD FUNCTION

For a reinforcement-learning extension, the reward can be defined as:

$$R_t = \lambda_1 Privacy + \lambda_2 Performance + \lambda_3 Interoperability - \lambda_4 Latency - \lambda_5 Communication$$

A large penalty should be applied if a mandatory privacy requirement is violated:

$$R_t = -\infty$$

for:

$$Privacy < Privacy_{minimum}$$

This prevents the model from selecting a high-performance but privacy-inadequate mechanism.

XXVII. DISCUSSION

The proposed framework addresses a fundamental limitation of static cryptographic architectures. Cryptographic mechanisms are often selected according to broad ecosystem preferences, but identity transactions are heterogeneous.

For example, an academic certificate presentation may require only selective disclosure. In contrast, a financial eligibility application may require proving several predicates simultaneously. A mobile wallet may prioritize low computational cost, while a high-assurance government application may prioritize privacy and expressive proof capabilities.

The adaptive approach therefore introduces a separation between:

1. **Identity representation**
2. **Credential representation**
3. **Verification policy**
4. **Cryptographic proof mechanism**

This separation allows the same identity ecosystem to accommodate multiple cryptographic technologies.

The approach is consistent with the modular nature of W3C DID and VC specifications. DID Core does not require a particular underlying technology, while the VC Data Model provides a general credential and presentation model that can accommodate multiple securing mechanisms. (W3C)

XXVIII. SECURITY–PERFORMANCE TRADE-OFF

A critical observation is that privacy and performance cannot always be optimized simultaneously.

A simplified optimization can be written as:

$$U = \lambda P + (1 - \lambda)Perf$$

where:

$$0 \leq \lambda \leq 1$$

When:

$$\lambda \to 1$$

privacy dominates.

When:

$$\lambda \to 0$$

performance dominates.

However, the proposed framework recommends that mandatory privacy constraints should be enforced before utility optimization.

This distinction is important because a privacy violation should not be treated as merely another small performance penalty.

XXIX. INTEROPERABILITY CONSIDERATIONS

Interoperability is critical for practical decentralized identity deployment.

The proposed framework should therefore avoid modifying the underlying DID and VC data models unnecessarily.

XXX. LIMITATIONS

The proposed framework has several limitations.

1. Cryptographic Benchmark Dependence

Performance varies significantly according to hardware, software libraries, parameter sizes, and implementation quality.

2. Standardization Evolution

DID and VC standards are evolving. For example, DID Core 1.1 is currently a W3C Candidate Recommendation Snapshot rather than the final W3C Recommendation [8]. (W3C)

3. Privacy Cannot Be Reduced to One Number

A single privacy score is useful for decision-making but cannot fully capture all privacy risks.

4. Metadata Leakage

Even a cryptographically private proof can leak information through network-level metadata, timing, endpoint identifiers, or repeated interaction patterns.

5. Trust Governance

Cryptographic verification does not automatically establish whether an issuer is trustworthy. The W3C VC model explicitly leaves issuer trust determination to the verifier and ecosystem governance [2]. (W3C)

6. Need for Empirical Validation

The mathematical selection framework requires implementation and benchmarking before quantitative performance claims can be made.

XXXI. FUTURE RESEARCH

Future research can extend AZSF in several directions.

A. AI-Based Selection

A machine-learning model can predict the optimal proof mechanism based on historical performance and transaction context.

B. Energy-Aware Selection

For mobile and IoT environments:

$$Score = Privacy + Performance + Interoperability - Energy$$

C. Quantum-Resistant Proof Selection

Future versions can incorporate post-quantum cryptographic primitives.

D. Cross-Chain Identity

The framework can be evaluated across multiple DID methods and distributed ledger environments.

E. Formal Verification

Formal methods can be used to verify policy enforcement and proof-selection logic.

F. Privacy Leakage Measurement

Future research should develop quantitative models for measuring correlation and metadata leakage.

G. Real-World Educational Identity

The framework can be implemented for university-issued credentials, including:

- degree certificates; student identity; scholarships; examination credentials; research credentials; employment verification.

XXXII. EXPECTED RESEARCH OUTCOMES

The proposed research is expected to establish:

1. A formal model for adaptive cryptographic proof selection.
2. A reference architecture for DID/VC-based privacy-preserving identity.

3. A policy language for specifying privacy and performance requirements.
4. A benchmark methodology for candidate ZKP mechanisms.
5. An implementation framework suitable for mobile and web wallets.
6. A foundation for AI-assisted cryptographic selection.
7. A privacy-preserving credential verification architecture suitable for education, e-governance, finance, and enterprise applications.

These are **research objectives and expected outcomes**, not reported experimental findings.

XXXIII. CONCLUSION

This paper proposed an **Adaptive ZKP Selection Framework (AZSF)** for privacy-preserving decentralized identity using Decentralized Identifiers and Verifiable Credentials. Existing decentralized identity systems provide an important foundation for user-controlled digital identity, but cryptographic proof mechanisms differ substantially in privacy capabilities, computational cost, communication requirements, interoperability, and expressiveness.

The proposed framework addresses this heterogeneity by introducing an adaptive decision layer that evaluates verification requirements before selecting an appropriate proof mechanism. The framework combines hard security constraints with multi-criteria optimization, thereby preventing cryptographic mechanisms that fail mandatory privacy requirements from being selected solely because of superior performance.

The proposed architecture supports selective disclosure, predicate verification, unlinkability, holder binding, replay resistance, and context-aware proof generation. Candidate mechanisms such as SD-JWT, BBS-based proofs, and general-purpose ZKP systems can be evaluated within a common framework.

The central contribution is therefore not the introduction of another standalone ZKP scheme, but the development of a **decision framework for selecting among cryptographic mechanisms according to the requirements of a particular identity transaction**.

Future implementation and empirical benchmarking will be necessary to validate the proposed scoring model, measure real-world performance, quantify privacy leakage, and evaluate the effectiveness of AI-assisted adaptive selection.

The proposed approach provides a pathway toward identity systems in which **privacy is treated as an explicit engineering objective rather than an incidental property of cryptographic implementation**.

REFERENCES

[1] M. Sporny, A. Guy, M. Sabadello, D. Longley, O. Steele, and C. Allen, “Decentralized Identifiers (DIDs) v1.0,” W3C Recommendation, Jul. 2022. W3C DID Core Recommendation

[2] M. Sporny, T. Thibodeau Jr., I. Herman, G. Cohen, and M. B. Jones, “Verifiable Credentials Data Model v2.0,” W3C Recommendation, May 2025. W3C Verifiable Credentials Data Model 2.0

[3] National Institute of Standards and Technology, “Privacy-Enhancing Cryptography: Zero-Knowledge Proof,” NIST Computer Security Resource Center. NIST Privacy-Enhancing Cryptography / ZKP

[4] T. Fett, K. Kuesters, and G. P. M. R. others, “Selective Disclosure for JSON Web Tokens (SD-JWT),” RFC 9901, IETF, Nov. 2025. RFC 9901: Selective Disclosure for JSON Web Tokens

[5] W3C Verifiable Credentials Working Group, “Data Integrity BBS Cryptosuites v1.0,” W3C Technical Specification. W3C Data Integrity BBS Cryptosuites

[6] "On cryptographic mechanisms for the selective disclosure of verifiable credentials," *Journal of Information Security and Applications*, vol. 83, 103789, Jun. 2024, doi: 10.1016/j.jisa.2024.103789. (ScienceDirect)

[7] "Analyzing and comparing the security of self-sovereign identity management systems through threat modeling," *International Journal of Information Security*, Springer, 2023. (Springer)

[8] M. Sporny and D. Zagidulin, "Decentralized Identifiers (DIDs) v1.1," W3C Candidate Recommendation Snapshot, Mar. 2026. W3C DID Core 1.1 Candidate Recommendation

[9] 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. (ScienceDirect)

[10] J. Sedlmeir, R. Smethurst, A. Rieger, and G. Fridgen, "Digital Identities and Verifiable Credentials," *Business & Information Systems Engineering*, vol. 63, pp. 603–613, 2021, doi: 10.1007/s12599-021-00722-y.

[11] L. D. Anh and J. H. Park, "Privacy-Preserving Identity Management System on Blockchain Using zk-SNARK," *IEEE Access*, 2023, doi: 10.1109/ACCESS.2022.3233828. (ResearchGate)

[12] A. Tobin and D. Reed, "The Inevitable Rise of Self-Sovereign Identity," Sovrin Foundation, 2017. Sovrin SSI White Paper

[13] W3C, "Verifiable Credentials Overview," W3C Technical Report, 2026. W3C Verifiable Credentials Overview

[14] W3C, "BBS Cryptosuite v2023," W3C Working Draft, 2023. The specification describes selective disclosure and unlinkable derived proofs using BBS signatures. (W3C)

[15] R. Chairattana-Apirom, D. Hofheinz, and S. Tessaro, "Tight Security for BBS Signatures," 2026. The work studies concrete security of BBS signatures and their relevance to privacy-preserving authentication and standardization. (arXiv)