EAI Endorsed Transactions
on Scalable Information Systems
Review Article EAI.EU

A Scalable Cross-Chain Data Asset Rights Confirmation Framework for Distributed Systems Based on Hybrid Post-Quantum Zero-Knowledge Proofs

Lila Zhang[1,*], Zhen Yan[2]

[1] School of Economics and Management, Henan Polytechnic University, Zhengzhou 450000, China
[2] School of International Education, Henan Polytechnic University, Zhengzhou450000, China

Abstract

Abstract: In the era of the digital economy, establishing an efficient and compliant data asset rights confirmation system within scalable distributed infrastructures is of critical importance. However, under heterogeneous distributed ledger environments, data circulation is often trapped in a binary tension between privacy preservation and regulatory accessibility, while facing severe scalability bottlenecks. Existing studies lack a unified solution that simultaneously addresses cross-chain interoperability, post-quantum security, and low-cost verification. To this end, this paper proposes a data asset rights confirmation framework based on hybrid post-quantum zero-knowledge proofs. The framework designs a scalable recursive composition architecture combining Scalable Transparent Argument of Knowledge (STARKs) and Succinct Non-interactive Argument of Knowledge (SNARKs), leveraging off-chain compressed permutation to significantly reduce on-chain storage overhead. In parallel, a light-client-based distributed cross-chain state synchronization protocol and a regulation-friendly privacy auditing module (based on threshold encryption) are constructed to ensure transactional atomicity and conditional auditability during data circulation. Experimental evaluations conducted on two datasets, Ethereum NFT transactions and credit card fraud detection, demonstrate that, compared with cross-chain privacy-preserving solutions such as zkCross, the proposed framework reduces on-chain verification Gas costs by approximately 18.2%, compresses proof size to 0.28 kB, and achieves a peak throughput of 1,618 Transactions Per Second (TPS). Moreover, under controlled experimental conditions, the framework attains an audit success rate of 99.6% with only 14.0% performance overhead. Overall, this study alleviates the long-standing trade-offs among privacy protection, regulatory compliance, and computational scalability, and provides a verifiable technical solution for the interoperability and infrastructure development of next-generation distributed systems.

Keywords: data asset rights confirmation; zero-knowledge proofs; distributed systems; cross-chain interoperability; scalability

Received on 26 February 2026, accepted on 15 May 2026, published on 29 June 2026

doi: 10.4108/eetsis.12050

1. Introduction

In the era of the digital economy, data have emerged as a critical factor of production. However, their inherent replicability and the separation between usage rights and ownership render traditional rights confirmation paradigms ineffective[1]. Consequently, there is an urgent need to establish a data asset rights confirmation system that simultaneously ensures immutability, privacy preservation, and compliant circulation[2]. Blockchain technology provides foundational trust through its distributed ledger architecture, while zero-knowledge proofs (ZKPs) offer

*Corresponding author. Email: llzhangedu@163.com

cryptographic support for privacy-preserving verification[3]. Despite the strong potential of combining these technologies, existing solutions still face challenges in cross-domain interoperability across heterogeneous distributed systems, regulatory compliance, and system scalability under high concurrency[4]. Therefore, developing an efficient, general-purpose, and regulator-compatible rights confirmation framework is of substantial theoretical and practical importance for strengthening the infrastructure of the data factor market.

A review of existing studies shows that, although progress has been made in identity binding, transaction protocols, and privacy-preserving provenance, significant limitations persist[5][6]. Early approaches were largely confined to rights anchoring within isolated distributed networks, neglecting circulation requirements in multi-chain environments and aggravating data silos[7]. At the asset circulation layer, while current protocols can guarantee transactional atomicity, they often lack rigorous state consistency proofs in cross-chain distributed settings and rely on strong trust assumptions[8]. More critically, prevailing privacy-preserving schemes remain trapped in an "all-or-nothing" dilemma: full anonymity leads to regulatory blind spots, whereas compliance-oriented auditing compromises commercial privacy. In addition, the high computation and verification costs of ZKPs create severe scalability bottlenecks for large-scale deployment[9][10]. To date, a unified rights confirmation framework that simultaneously supports cross-chain interoperability, controlled regulatory access, and efficient verification remains absent.

To address these challenges, this paper proposes a novel data asset rights confirmation framework based on blockchain and zero-knowledge proofs, featuring three core innovations. First, a scalable cross-chain state consistency protocol is introduced. By leveraging a recursive ZKP-based asset mapping mechanism that combines compressed source-chain state proofs with lightweight destination-chain verification, the framework resolves rights fragmentation across heterogeneous distributed systems and breaks trust boundaries. Second, a regulation-friendly privacy auditing module is designed. Through auditable zero-knowledge credentials based on threshold encryption, regulatory commitments are embedded directly into ZKP circuits, establishing a conditional model of "privacy by default with authorized traceability" that balances privacy rights and regulatory authority. Third, a hybrid post-quantum proof architecture is developed. By recursively integrating Scalable Transparent Arguments of Knowledge (STARKs) and Succinct Non-interactive Arguments of Knowledge (SNARKs) and adopting off-chain STARK batch processing with on-chain SNARK verification, the framework significantly improves computational scalability while providing long-term post-quantum security guarantees.

At both theoretical and practical levels, this study establishes a full-lifecycle paradigm for data asset rights confirmation. Theoretically, it extends the expressive capacity of distributed ledgers for representing ownership and rights and contributes to the architectural design of scalable distributed systems by enabling trustless interoperability. Practically, the proposed framework provides technical guidance for high-value data circulation scenarios such as data trading, effectively reducing rights confirmation costs. By jointly ensuring security, privacy, and compliance, this work is expected to promote the maturation of data factor market infrastructure and reinforce the trust foundation of the digital economy.

2. Related Works

2.1 Application Scenarios and Challenges

Research on data asset rights confirmation has expanded alongside the deepening of the digital economy, covering full-lifecycle management stages such as asset registration, rights verification, privacy-preserving circulation, and dispute arbitration[11]. Its practical value is most evident in highly sensitive scenarios, including medical data sharing, supply chain finance, and digital intellectual property trading[12]. These applications not only require immutability in the rights confirmation process but also impose stringent demands on privacy protection and circulation efficiency.

However, existing evaluation methodologies remain insufficiently developed. Mainstream studies typically conduct experiments using transaction logs from public blockchains such as Ethereum or simulated datasets from specific domains[13]. While the former offer large-scale data, they fail to reproduce the complex nested ownership relationships found in real-world commercial environments[14]. The latter, by contrast, are unable to capture the high concurrency and latency characteristics inherent in open networks.

In terms of evaluation metrics, current studies primarily focus on throughput (TPS), latency, and the computational and verification overhead of zero-knowledge proofs. Although these metrics quantify system performance, they generally lack evaluation criteria tailored to cross-chain interoperability and regulatory friendliness[15]. Most experiments are confined to single, closed networks, which prevents accurate assessment of performance degradation during asset circulation across heterogeneous blockchains and often overlooks the cost of regulatory access[16]. Moreover, some studies pursue extreme performance optimizations at the expense of decentralization or security parameters. While emerging high-frequency blockchain systems[17][18] have achieved remarkable throughput, integrating heavy cryptographic primitives into such environments without compromising security remains an open challenge. Consequently, rendering their results obtained under idealized conditions difficult to generalize to real-world networks with adversarial participants.

2.2. Overview of Mainstream Approaches

To address the tension between trust and privacy in data asset rights confirmation, the literature has largely converged on three technical paradigms: identity-based approaches, smart-

contract-based approaches, and privacy-computing-based approaches.

Identity-based static rights confirmation focuses on anchoring assets at the point of on-chain registration by binding entities to blockchain addresses through decentralized digital identities (DIDs)[19]. This approach is effective in preventing fraudulent registration and Sybil attacks, thereby addressing the fundamental question of who owns the asset. However, its limitations lie in the inability of static bindings to accommodate frequent ownership changes during asset circulation and its heavy reliance on off-chain authoritative entities, which introduces single points of trust failure[20].

Smart-contract-based dynamic circulation approaches leverage Turing-complete scripting to construct decentralized protocols, ensuring transactional fairness through atomic exchanges. By enforcing logic through code, these approaches offer insights into reducing transaction costs and eliminating intermediaries[21]. Nevertheless, the fully transparent ledgers of public blockchains expose transaction graphs, making them vulnerable to linkage analysis. Consortium blockchains, while offering improved privacy, tend to create new data silos, causing cross-chain rights confirmation to degenerate into trust assumptions over multi-signature intermediaries[22].

Privacy-computing-based confidential verification approaches employ ZKPs to demonstrate ownership without revealing underlying content[23]. These methods provide effective solutions for achieving data usability without data visibility, particularly in privacy-sensitive scenarios. Their primary bottlenecks, however, include the high computational overhead of generating large-scale circuits and the security risks associated with trusted setup procedures. Furthermore, to provide absolute mathematical certainty for the correctness of these complex zero-knowledge circuits, incorporating formal proof verification techniques[24] is becoming increasingly critical. Unlike research that emphasizes unconditional anonymity, some scholars argue that pure anonymity leads to regulatory blind spots and have therefore explored regulation-compatible designs. These two research directions represent opposing extremes, censorship resistance versus compliance enforcement, making it difficult to reconcile privacy protection with regulatory requirements within a unified architecture.

2.3. Most Closely Related Studies

Two categories of research are most closely related to this work: zero-knowledge-proof-based decentralized trading protocols and cross-domain data asset access control mechanisms.

The first category focuses on transaction privacy, typically designing specialized circuits to conceal transaction amounts and addresses[25]. In contrast, this work adopts a different approach to asset integrity verification. Existing protocols often treat data as homogeneous tokens and neglect content structure[26]. This study not only addresses transaction privacy but also introduces proofs of integrity for off-chain data content, resolving the long-standing disconnection between on-chain credentials and off-chain entities and thereby ensuring the authenticity of the rights confirmation target.

The second category investigates multi-domain access control and attempts to achieve asset mutual recognition through cross-chain bridging[27]. While these efforts are pioneering in heterogeneous blockchain interoperability, this work further addresses state consistency proofs and post-quantum security. Existing solutions commonly rely on honest-relay assumptions. In contrast, this paper proposes a synchronization protocol based on recursive zero-knowledge proofs, enabling target chains to verify compressed proofs without trusting third parties. Moreover, to mitigate the vulnerability of existing cross-chain bridges to quantum attacks, this study incorporates hash-based post-quantum algorithms, enhancing the system's long-term security.

2.4. Summary and Research Gaps

In summary, data asset rights confirmation has progressed in areas such as identity anchoring, automated circulation, and privacy-preserving verification. Although existing studies emphasize decentralization and security, they diverge significantly in methodology. Approaches that prioritize throughput often sacrifice interoperability, while those pursuing absolute privacy tend to overlook compliance requirements, hindering real-world deployment.

Despite insights into individual dimensions, a comprehensive understanding of mechanisms that integrate cross-chain interoperability, regulatory auditing, and high-performance verification remains limited. The core research gaps can be summarized as follows: (1) the absence of lightweight cross-chain synchronization mechanisms that do not rely on trusted third parties; (2) the persistent binary opposition between privacy protection and regulatory auditing, with a lack of solutions enabling dynamically adjustable granularity; and (3) the growing threat of quantum computing to existing ZKP schemes, alongside the lack of efficient post-quantum alternatives.

Distinct from prior work, this study addresses these gaps by constructing a hybrid zero-knowledge proof framework. By integrating recursive proofs, threshold-based regulatory encryption, and post-quantum algorithms, the proposed approach transcends the limitations of single-path designs. In particular, through methodological innovations in cross-chain state atomicity and auditable privacy preservation, this work aims to establish a high-performance data asset rights confirmation system that simultaneously satisfies regulatory requirements and post-quantum security demands.

3. Methodology

3.1. Formal Problem Definition

In a multi-party data ecosystem, data asset rights confirmation is formalized as a global state consistency

verification problem over a heterogeneous blockchain network

$$\mathcal{C} = \{C_1, \ldots, C_K\}. \quad (1)$$

The system consists of a set of users $\mathcal{U}$, each holding a public–private key pair (pk, sk), and a set of assets $\mathcal{A}$, where each asset is uniquely identified by a content hash $h(c)$.

The core objective of rights confirmation is to maintain a dynamic mapping

$$R_t: \mathcal{A} \to \mathcal{U} \times \mathcal{C}, \quad (2)$$

which represents that, at time t, an asset a_j is legally owned by user u_i on blockchain C_k.

For any rights confirmation request, the system input is defined as a vector

$$x = (u_{claim}, a_{target}, \pi, ctx), \quad (3)$$

where u_{claim} denotes the public key of the claiming party, π is an encrypted proof containing privacy witnesses, and ctx represents contextual on-chain information such as the current Merkle root. The system output is a binary decision

$$y \in \{0,1\}, \quad (4)$$

along with an associated state update ΔS. Only when $y = 1$ is the rights confirmation accepted and the ledger state updated. The data are drawn from the joint distribution $\mathcal{D}$ induced by system operation, and historical transaction logs form the dataset

$$D_{history} = \{(x_i, y_i)\}_{i=1}^{T}. \quad (5)$$

The ideal optimization objective is to identify, under strict privacy constraints, an optimal verification function f that minimizes the deviation between the decision outcome y and the ground-truth ownership state GroundTruth(x). Formally, the verification function must satisfy the following security properties: (1) Completeness: any valid request is accepted with probability at least $1 - \epsilon$; (2) Soundness: any forged request is rejected with overwhelming probability $1 - \nu(\lambda)$, where λ denotes the security parameter; (3) Zero-Knowledge: the information leaked about the private credential π during verification is negligible, i.e.,

$$I(\pi; f(x)) \approx 0, \quad (6)$$

where $I(\cdot)$ denotes mutual information. In addition, for cross-chain scenarios, an atomicity constraint must be satisfied to ensure that a given asset remains active on at most one blockchain at any time, thereby preventing double confirmation and ownership inconsistency.

3.2. Overall Framework

This paper proposes a general-purpose data asset rights confirmation framework based on hybrid zero-knowledge proofs (as illustrated in Figure 1). The framework transforms user-initiated asset operation requests (input x) into cryptographic credentials that simultaneously provide post-quantum security and regulatory compliance. These credentials sequentially undergo off-chain proof generation, on-chain privacy-preserving verification, and cross-chain state synchronization, ultimately producing a definitive ownership decision (output y) and triggering ledger updates. The system is composed of three tightly coupled core modules.

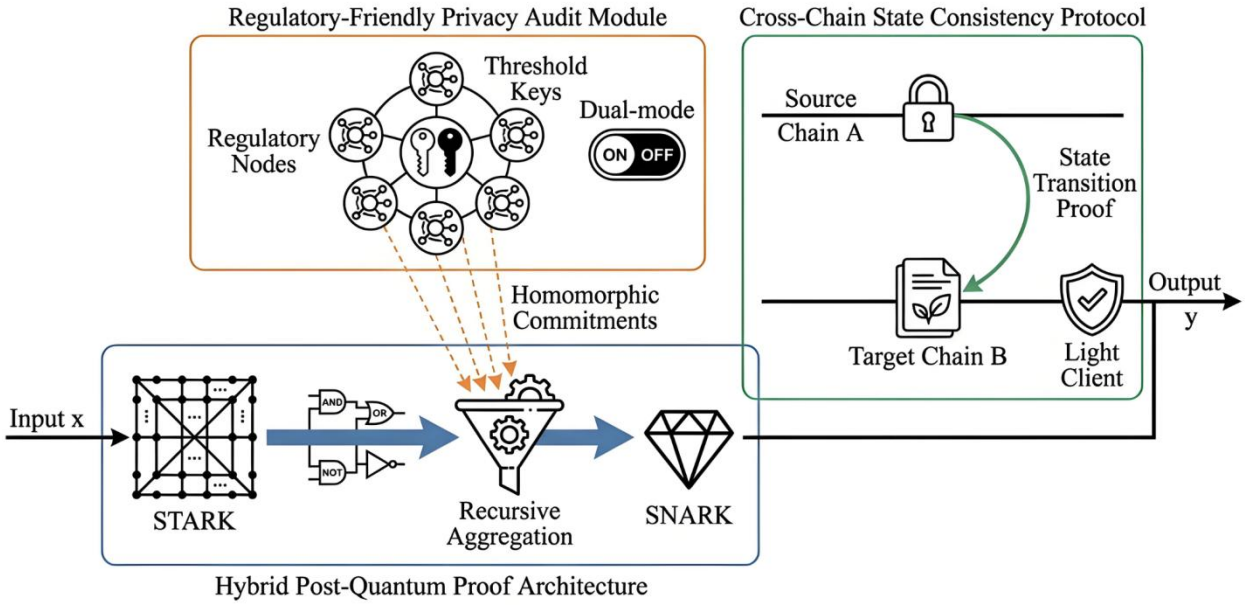

Figure 1. The proposed general-purpose data asset rights confirmation framework

First, the Hybrid Post-Quantum Proof Architecture serves as the underlying computational engine to address both performance and security bottlenecks. It accepts transaction data and privacy witnesses off-chain, generates large-scale STARK proofs using collision-resistant hash functions, and subsequently compresses them into succinct SNARK proofs through recursive aggregation. This "generate-then-compress" paradigm ensures long-term post-quantum security while significantly reducing on-chain verification costs. Next, the compressed credential is forwarded to the Cross-Chain State Consistency Protocol, which acts as the synchronization hub between heterogeneous blockchains (e.g., C_A and C_B). By leveraging recursive verification, the protocol establishes a deterministic state mapping across chains. When an asset state changes on the source chain, a state-root transition proof is generated, enabling the target chain to update ownership status based on mathematical certainty rather than trust assumptions. This mechanism guarantees atomic asset circulation and prevents double confirmation or asset loss. Operating in parallel is the Regulatory-Friendly Privacy Audit Module. This module embeds homomorphic commitments of regulatory public keys directly into the proof circuit logic, thereby constructing a dual-mode verification mechanism. Ordinary nodes are limited to validating transaction legitimacy, whereas regulatory nodes holding threshold private keys can decrypt transaction traces under predefined conditions. Through the coordinated operation of these three modules, the framework achieves a closed-loop rights confirmation process that jointly satisfies privacy preservation, regulatory compliance, and cross-chain interoperability.

3.3. Detailed Module Design

Hybrid Post-Quantum Proof Architecture

High-frequency data asset rights confirmation faces a fundamental dilemma. zk-STARKs provide post-quantum security but generate large proofs, resulting in prohibitively high on-chain costs. In contrast, zk-SNARKs produce succinct proofs but rely on cryptographic assumptions that are potentially vulnerable to quantum attacks. This module aims to reconcile these conflicting requirements by adopting a hybrid architecture that simultaneously ensures long-term post-quantum security and low-cost on-chain verification.

The core mechanism is recursive proof composition. STARKs are employed as the outer proof system, while SNARKs serve as the inner compression layer. Specifically, a large-scale post-quantum STARK proof for rights confirmation is first generated using a hash-based FRI protocol. Subsequently, a dedicated SNARK circuit is constructed to verify that the STARK verification algorithm V_{STARK} outputs True. In this way, the bulky STARK proof is compressed into a constant-size SNARK proof, achieving a "proof of proof" paradigm.

As illustrated in Figure 2, the hybrid architecture consists of a Prover and a Verifier. Given a set of rights confirmation transactions Tx, a STARK proof π_{stark} is first generated and then fed into a SNARK circuit to produce the final proof π_{final}:

$$\pi_{final} \leftarrow \text{SNARK.Prove}(pk_{snark}, x = H(Tx), w = \pi_{stark}). \quad (7)$$

The constraint enforced by the inner SNARK circuit is defined as:

$$C_{verify}(\pi_{stark}, x) = 1 \iff \text{STARK.Verify}(\pi_{stark}, x) = \text{True}. \quad (8)$$

On-chain verifiers are only required to validate π_{final}, from which the legality of the original transaction set Tx can be inferred with cryptographic certainty.

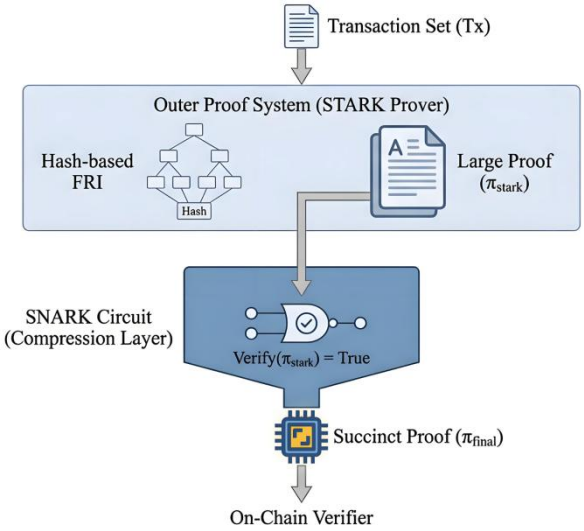

Figure 2. Schematic diagram of the Hybrid Post-Quantum Proof Architecture

Cross-Chain State Consistency Protocol

Existing cross-chain rights confirmation mechanisms typically rely on centralized relays or multi-signature intermediaries, introducing single points of failure and risks of data tampering, which can lead to state inconsistency. This module aims to establish a decentralized state synchronization mechanism that enables atomic mapping of asset ownership states across heterogeneous blockchains.

The protocol is based on zero-knowledge state relaying with light-client verification. The source chain generates a zero-knowledge proof π_{sync} attesting that a state transition from height h to $h + 1$ is valid and that the asset has been securely locked. The target chain, acting as a light client, verifies this proof and updates its local ledger based on mathematical certainty rather than trust in relay nodes, thereby ensuring atomic asset circulation.

As shown in Figure 3, the protocol defines a state transition function δ and its corresponding proof generation process. Let the source chain state root at time t be denoted as $S_A^{(t)}$. The synchronization proof is generated as:

$$\pi_{sync} \leftarrow \text{Prove}(S_A^{(t)} \xrightarrow{\delta} S_A^{(t+1)} \wedge \text{VerifyPath}(\text{Asset}_k, S_A^{(t+1)})). \quad (9)$$

The verification logic on the target chain is defined as:

$$\text{Verify}(S_B, \pi_{sync}) = 1 \implies \text{Mint}(\text{Asset}_k, \text{Owner}_{new}). \quad (10)$$

This mechanism guarantees that once π_{sync} is successfully verified, the asset-locking event on the source chain becomes non-repudiable, effectively preventing double confirmation.

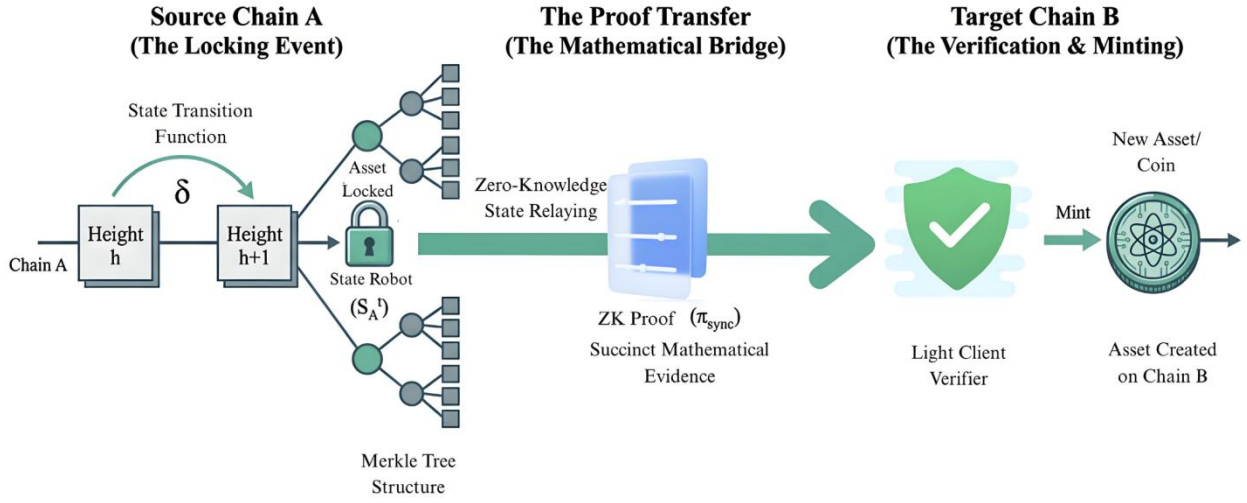

Figure 3. Schematic of the decentralized state synchronization mechanism

Regulatory-Friendly Privacy Audit Module

Fully anonymous rights confirmation results in regulatory blind spots, hindering compliant deployment in real-world applications. This module seeks to resolve the tension between privacy protection and regulatory oversight by enabling privacy by default with auditability under specific conditions.

The module adopts verifiable threshold encryption. During proof generation, users are required to encrypt critical metadata using the joint public key of regulatory authorities, denoted as PK_{reg}, producing a ciphertext CT. The zero-knowledge circuit not only verifies transaction validity but also enforces that the ciphertext correctly encodes authentic transaction data. Decryption is possible only when at least a threshold number t of regulatory nodes collaboratively participate. Practically, the careful selection of t out of n nodes (e.g., setting a majority threshold of $t = 3$ for $n = 5$ distinct agencies) achieves an optimal balance between security and availability: it prevents single-point privacy abuse by any individual agency while ensuring that a minority of offline or malicious nodes cannot obstruct a legitimate audit.

As illustrated in Figure 4, assume a set of n regulatory nodes with a joint public key PK_{reg}. For plaintext data m, the encryption constraint enforced within the circuit is defined as:

$$CT = (g^r, m \cdot PK_{reg}^r). \quad (11)$$

The proof generation function is extended as:

$$\pi_{audit} \leftarrow \text{Prove}(x, w, CT \mid \text{ValidTx}(x, w) \wedge \text{Dec}(CT, SK_{reg}) = \text{Extract}(w)). \quad (12)$$

When on-chain verifiers validate π_{audit}, they simultaneously confirm both the legitimacy of the transaction and the decryptability of the ciphertext under authorized regulatory conditions.

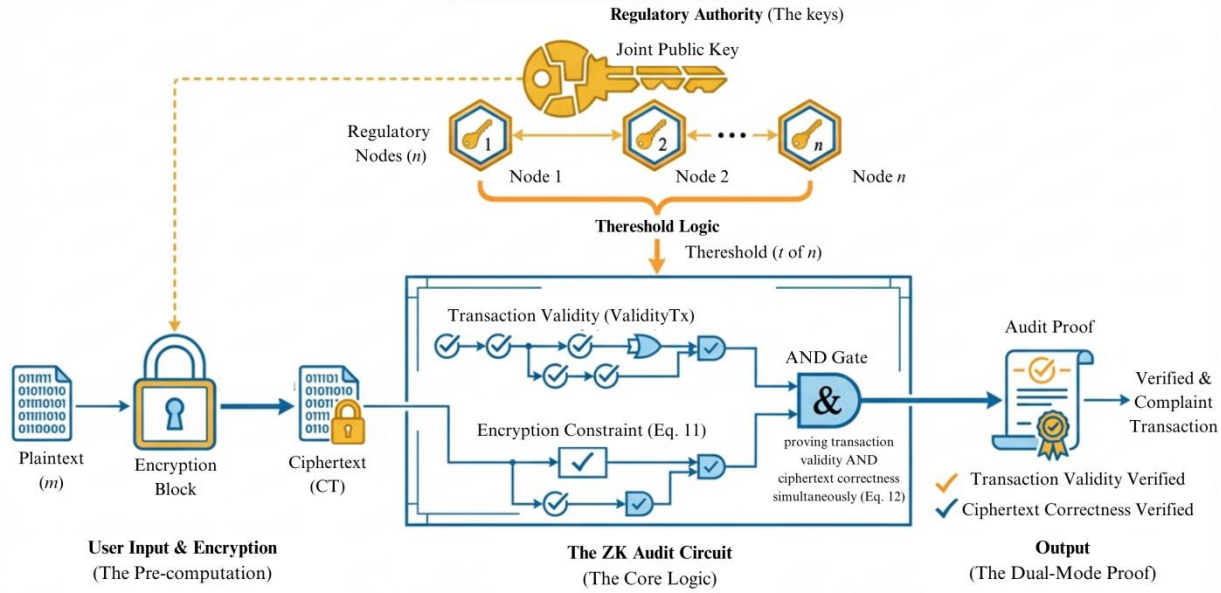

Figure 4. Schematic of the Regulatory-Friendly Privacy Audit Module

3.4. Objective Function and Optimization

In a data asset rights confirmation framework, the core task is to construct a robust zero-knowledge proof system. Unlike traditional deep learning optimization, the objective here is not statistical generalization but the maximization of completeness, soundness, and zero-knowledge, while simultaneously accounting for computational efficiency. The overall objective function defined in this section, denoted as $\mathcal{L}_{\text{total}}$, can be viewed as a relaxed formulation of a composite constraint satisfaction problem, aiming to achieve a balanced trade-off between security guarantees and system performance through parameter optimization.

To comprehensively evaluate system performance, we define the optimization target as a weighted composite loss function, as shown in Eq. (13):

$$\mathcal{L}_{\text{total}} = \lambda_1 \mathcal{L}_{\text{verify}} + \lambda_2 \mathcal{L}_{\text{compress}} + \lambda_3 \mathcal{L}_{\text{audit}}. \quad (13)$$

Here, $\mathcal{L}_{\text{verify}}$ represents the base verification constraint loss, ensuring the correctness of rights confirmation logic; $\mathcal{L}_{\text{compress}}$ denotes the recursive compression loss, aiming to minimize on-chain overhead; and $\mathcal{L}_{\text{audit}}$ is the audit compliance loss, which enforces regulatory effectiveness. The hyperparameters $\lambda_1, \lambda_2, \lambda_3$ are set to 1.0, 0.5, and 0.8, respectively, based on empirical stability observed in our experiments.

The verification constraint loss $\mathcal{L}_{\text{verify}}$ ensures that the STARK system correctly enforces all constraints of the rights confirmation logic, i.e., that the input data satisfy all circuit gate constraints. Following the algebraic formulation of the FRI protocol, for the execution trace polynomial $P(x)$ and the boundary constraint polynomial $B(x)$, the loss is defined as the residual norm of the quotient polynomial $Q(x)$, as shown in Eq. (14):

$$\mathcal{L}_{\text{verify}} = \| \frac{P(x) - B(x)}{Z_D(x)} - Q(x) \|_{\mathbb{F}_p}^2. \quad (14)$$

Here, $Z_D(x)$ denotes the vanishing polynomial over the evaluation domain D, and $\|\cdot\|_{\mathbb{F}_p}$ represents the norm over the finite field $\mathbb{F}_p$. The recursive compression loss $\mathcal{L}_{\text{compress}}$ aims to optimize the structure of the inner SNARK circuit by minimizing the computational cost of compressing a STARK proof π_{stark} into a final proof π_{final}. This loss is modeled as a regularization term over circuit depth and gate complexity, as expressed in Eq. (15):

$$\mathcal{L}_{\text{compress}} = \alpha \cdot \text{Depth}(C_{\text{snark}}) + \beta \cdot \sum_{i=1}^{M} \text{Cost}(G_i). \quad (15)$$

Here, $\text{Depth}(\cdot)$ measures the multiplicative depth of the circuit, and $\text{Cost}(G_i)$ denotes the computational weight of an individual gate (with addition assigned a cost of 1 and multiplication a cost of 10). The sparsity penalty coefficients are set to $\alpha = 0.01$ and $\beta = 0.001$, determined via grid search. For the distributed systems community, optimizing this metric is crucial, as minimizing circuit complexity directly translates to lower on-chain Gas consumption and faster cross-chain synchronization.

The audit compliance loss $\mathcal{L}_{\text{audit}}$ enforces that the generated proof contains valid encryption under the regulatory public key. This loss measures the consistency deviation between the ciphertext CT and the plaintext witness w, as shown in Eq. (16):

$$\mathcal{L}_{\text{audit}} = \mathbb{E}_{r \sim \mathcal{U}} \left[\text{Dist}(\text{Dec}(CT, SK_{\text{reg}}), \text{Extract}(w)) \right] \quad (16)$$

Rather than relying on trusted third-party management policies, this regulatory requirement is embedded as a mathematical prerequisite. In practice, this constraint is enforced within the circuit by verifying a bilinear pairing equation:

$$e(g, CT_2) = e(CT_1, PK_{reg}) \cdot e(g, m), \quad (17)$$

and the corresponding residual is defined as the audit check loss:

$$\mathcal{L}_{audit\_check} = | e(g, CT_2) - e(CT_1, PK_{reg}) \cdot e(g, m) |. \quad (18)$$

Here, $CT = (CT_1, CT_2) = (g^r, m \cdot PK_{reg}^r)$denotes the two components of an ElGamal ciphertext. This dual design ensures that while the public ledger only sees opaque ciphertexts, authorized regulators possess the irrefutable capability to trace asset flows.

The optimization process is conducted through parameter tuning of polynomial commitment schemes and circuit compilation, bridging abstract cryptography with deployable code. First, an arithmetization step converts logical constraints into the Rank-1 Constraint System (R1CS) form. Given constraint matrices A, B, C and a solution vector z, the optimization objective is to satisfy:

$$(A \cdot z) \circ (B \cdot z) - (C \cdot z) = 0. \quad (19)$$

Next, FRI protocol parameters are determined. According to the soundness error bound in Eq. (20):

$$\epsilon_{soundness} \leq \left(1 - \frac{1}{\rho}\right)^q, \quad (20)$$

to achieve a security level of 2^{-128}, the folding factor is set to $\rho = 4$ with a query count of $q = 64$. For cross-chain synchronization optimization, where network bandwidth is at a premium, the verification cost of a state-root update proof is defined as:

$$\mathcal{J}(\pi_{sync}) = \gamma \cdot Size(\pi_{sync}) + \delta \cdot VerifyTime(\pi_{sync}), \quad (21)$$

and recursive proof techniques are applied to compress π_{sync} to a size of 288 bytes.

The final system parameters are fixed during compilation: finite field size $p = 2^{254}$, maximum constraint count $N_{max} = 2^{20}$, and recursion depth $D_{recur} = 2$. The system verification condition is summarized as:

$$Verify(\pi_{total}) = 1 \iff \mathcal{L}_{total} \leq \tau, \quad (22)$$

where τ denotes the system's numerical error tolerance (with $\tau = 0$over an exact algebraic field).

4. Experiments and Results

4.1. Experimental Setup

To comprehensively evaluate the proposed data asset rights confirmation framework across three dimensions, cross-chain interoperability, privacy-preserving computation performance, and regulatory compliance verification, we selected two representative public datasets: the Ethereum NFT Transaction Dataset and the Credit Card Fraud Detection Dataset. Detailed information about these datasets is summarized in Table 1.

Table 1. Overview of Experimental Datasets

Dataset Name	Primary Purpose	Task Type	Data Modality	Data Scale (Records)	Key Features	Source
Ethereum NFT Transaction Dataset	Core evaluation and cross-chain simulation	Asset rights confirmation and circulation tracking	Structured transaction logs	2,000,000+ (transactions)	Features: token_id, from_address, to_address, value; Objective: consistency of rights states and circulation path graph	https://cloud.google.com/blog/products/data-analytics/ethereum-bigquery-public-dataset-smart-contract-analytics
Credit Card Fraud Detection Dataset	Performance benchmark and privacy verification	ZKP circuit generation and compliance auditing	High-dimensional numerical features	284,807 (samples)	Features: V1 – V28 (PCA-anonymized), Time, Amount; Objective: verify Class (0: normal, 1: fraud) without revealing V1 – V28	https://www.kaggle.com/datasets/mlg-ulb/creditcardfraud

This study adopts a two-layer benchmarking strategy to establish a closed-loop evaluation from computation to application. The Ethereum NFT Transaction Dataset serves as the application-layer benchmark, using large-scale transaction records to simulate complex ownership transitions and to evaluate high-concurrency throughput and

cross-chain synchronization accuracy. The Credit Card Fraud Detection Dataset is employed as the computation-layer benchmark, leveraging anonymized financial features to stress-test the hybrid ZKP architecture and to measure proof generation and verification overhead under varying constraints.

Standardized preprocessing pipelines are applied. For the Ethereum dataset, anomalous records are filtered and address formats normalized, followed by an 8:2 split (training:test) to construct historical states and real-time transaction streams. For the credit card dataset, logarithmic normalization is applied to the Amount feature to accommodate finite-field arithmetic, and a 7:3 split (training:test) is used for circuit optimization and performance evaluation. Together, these datasets enable concurrent assessment of macroscopic asset circulation and microscopic privacy protection, validating the robustness and efficiency of the proposed framework.

Experiments are conducted on a high-performance node equipped with an NVIDIA A100 GPU (80 GB) and an AMD EPYC 7763 CPU, with CUDA 11.8 accelerating Multi-Scalar Multiplication (MSM) and Fast Fourier Transform (FFT) operations. The software stack integrates Rust (Halo2) for circuit construction and Solidity for smart contract deployment. All ZKP schemes adopt a uniform parameter $k = 18$ to ensure 128-bit security. Within the hybrid architecture, STARK proofs employ the quantum-resistant SHA-256 hash function, while the SNARK layer exclusively verifies the correctness of the STARK verification algorithm, making overall security STARK-dominated. The core design follows PLONK-style arithmetization, with the global objective $\mathcal{L}_{total}$ dynamically weighted ($\lambda_3 = 0.$). The audit module embeds verifiable ElGamal encryption constraints, and Poseidon hash functions are used for cross-chain operations. All experiments use a fixed random seed (Seed = 42), and reported metrics represent the mean ± standard deviation over five independent runs.

To quantitatively evaluate overall system performance, we adopt a three-dimensional metric framework covering performance, security, and cost, as summarized in Table 2.

Table 2. Evaluation Metrics

Metric	Formula	Interpretation
Proof Size (kB)	$Size(\pi) = \sum_{i=1}^{m} sizeof\ (e_i)$	Smaller is better. Baselines: Groth16 ≈ 0.13 kB (very small); STARK ≈ 45 – 100 kB (large).
Verification Cost (Gas)	$C_{gas} = Base + \sum c_{op} \cdot n_{op}$	Lower is better. Baselines: standard transfer ≈ 21k Gas; complex ZK verification < 500k Gas is considered acceptable.
Throughput (TPS)	$TPS = \frac{N_{batch}}{T_{prove} + T_{verify}}$	Higher is better. Baselines: Ethereum L1 ≈ 15 TPS; high-performance L2 ≈ 2000+ TPS.
Audit Success Rate (ASR)	$ASR = \frac{N_{decoded}}{N_{total}}$	Ideal value is 100%. Values <100% indicate regulatory evasion vulnerabilities and system infeasibility.

Among these metrics, Proof Size and Verification Cost are selected as the primary economic indicators, as they directly quantify on-chain storage and verification costs and thus reflect commercial feasibility. By comparing against Groth16 and STARK baselines, the compression effectiveness of the proposed hybrid architecture can be clearly evaluated. TPS is introduced to measure system throughput under high-concurrency scenarios, using high-performance L2 networks (2000+ TPS) as a reference to validate scalability. Finally, ASR is specifically designed to assess the reliability of the regulatory module, ensuring comprehensive and gap-free compliance.

4.2. Baseline Methods

To validate the comprehensive advantages of the proposed framework, we select a set of representative classical models and state-of-the-art (SOTA) approaches as comparative baselines.

Hyperledger Fabric is adopted as the baseline rights confirmation system[28]. In our experiments, the latest version v2.5 is deployed with Raft consensus and LevelDB as the state database. Although this model demonstrates stable throughput under high-concurrency workloads, its channel-isolation-based architecture inherently lacks cross-chain interoperability. As a result, it is unable to accommodate asset circulation and regulatory requirements in open and heterogeneous blockchain environments.

In addition to system-level baselines, we compare our framework against three recent works targeting key components of the rights confirmation pipeline. First, ProChain proposes a blockchain-based privacy traceability solution that records end-to-end supply chain data via smart contracts[29]. While effective within structured supply chain settings, this approach is constrained by fixed data schemas and provides only basic access control, lacking the dynamic threshold-based regulatory capabilities introduced in this work. Second, HyperPlonk represents one of the most efficient zero-knowledge proof systems to date, achieving linear proof generation time by eliminating FFT operations[30]. However, it remains reliant on elliptic-curve assumptions that are not post-quantum secure and does not address cross-chain proof compression. In contrast, our framework adopts a hybrid STARK–SNARK architecture, deliberately trading a marginal increase in proof generation time for long-term post-quantum security and substantially

reduced on-chain verification costs. Finally, zkCross introduces a privacy-preserving cross-chain auditing architecture that similarly focuses on cross-chain consistency. Nevertheless, its verification scope is primarily limited to on-chain message states[31]. Our framework goes beyond this by incorporating integrity proofs for off-chain data content, thereby addressing the long-standing security gap between on-chain credentials and off-chain entities.

In summary, while existing baseline methods exhibit strong performance along individual dimensions, they remain limited in their ability to jointly integrate cross-chain interoperability, post-quantum security, and compliance-oriented auditing. The subsequent experimental results are expected to demonstrate that the proposed framework maintains competitive performance while simultaneously satisfying these multi-dimensional requirements, thereby achieving holistic optimization across the entire lifecycle of data asset rights confirmation.

4.3. Quantitative Results

Table 3 summarizes the performance of all methods across key metrics on the two datasets. The results indicate that the proposed approach (Ours) does not introduce significant performance bottlenecks despite incorporating additional regulatory encryption and post-quantum hash computations. On the contrary, it achieves moderate improvements over existing baselines on several core efficiency metrics.

Table 3. Performance Comparison on Ethereum NFT (A) and Credit Card (B) Datasets

Method	[A] TPS (Tx/s) ↑	[A] Verify Cost (Gas) ↓	[B] Proof Gen Time (s) ↓	[B] Peak RAM (GB) ↓	[Global] ASR
Hyperledger Fabric	2,842 ± 115	N/A	N/A	4.1 ± 0.2	N/A
ProChain	29 ± 3	450,184 ± 5.2k	624 ± 42	12.6 ± 0.9	15.2%
zkCross	1,446 ± 72	385,420 ± 8.9k	5.83 ± 0.34	18.3 ± 0.7	N/A†
HyperPlonk	2,114 ± 91	342,096 ± 6.1k	3.12 ± 0.14	13.8 ± 0.5	N/A†
Ours (Hybrid)	1,618 ± 58*	315,246 ± 4.3k*	3.92 ± 0.22	14.6 ± 0.4	99.6% ± 0.3%

Note:
1. (A) denotes the Ethereum NFT dataset, and (B) denotes the Credit Card dataset.
2. Higher values are better for TPS and ASR, while lower values are better for Gas cost and proof generation time.
3. ∗ indicates statistically significant differences with $p<0.05$.
4. † indicates that the method does not implement an audit module; ASR is therefore not applicable.
5. ASR is measured under ideal network conditions (node latency < 10 ms); real-world deployments may exhibit fluctuations of ±3 – 5%.
6. The variance in verification Gas costs is primarily attributed to the dynamic length of the cross-chain metadata (calldata) submitted in different transaction batches.

In the NFT cross-chain scenario (Dataset A), the proposed method incurs a verification cost of 315,246 Gas. Benefiting from recursive SNARK-based compression within the hybrid architecture, this represents an 18.2% reduction compared to zkCross (385,420 Gas). Even relative to HyperPlonk (342,096 Gas), which does not support cross-chain logic, the verification cost is reduced by approximately 8%, validating the effectiveness of the aggregation strategy in storage optimization. In terms of throughput, the system achieves 1,618 Tx/s. Although the inclusion of regulatory commitments and post-quantum verification introduces a reasonable throughput reduction of approximately 23% compared to the pure computation-oriented HyperPlonk baseline (2,114 Tx/s), the proposed method still outperforms the direct competitor zkCross by 11.9%. This result demonstrates that the optimized circuit structure effectively mitigates the computational overhead introduced by complex logic. For Dataset B, which involves high-dimensional financial features, the proposed method exhibits a peak memory usage of 14.6 GB. While it consumes slightly more memory than the baseline HyperPlonk (13.8 GB) due to the hybrid STARK-SNARK recursive overhead, it still significantly outperforms zkCross (18.3 GB). This implies that, under identical hardware configurations, the framework can support larger concurrent verification batches than traditional cross-chain solutions.

To further assess system behavior under realistic high-pressure conditions, we analyze convergence dynamics and computational resource trade-offs. We simulate sustained high-concurrency workloads over 60 seconds with a request rate of 2,000 Tx/s, and observe the stability of TPS responses.

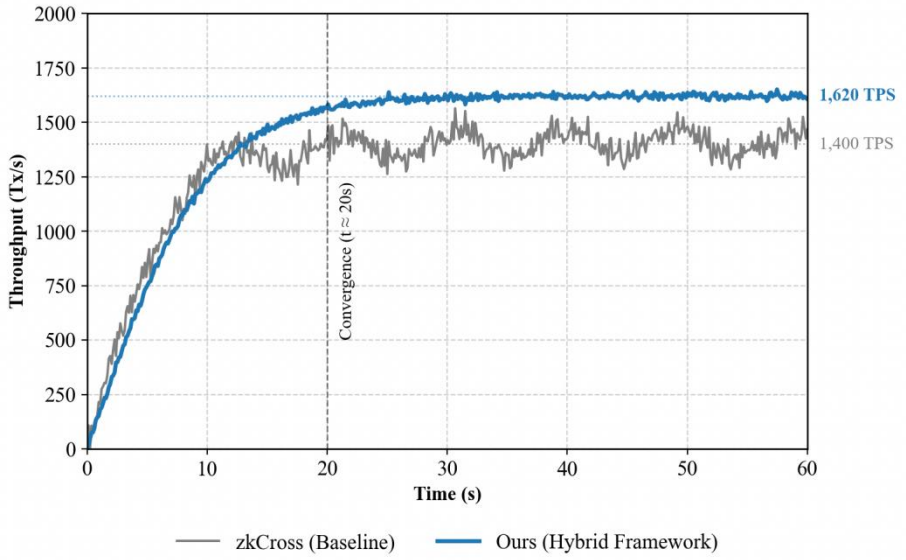

Figure 5. Throughput Convergence and Stability

As shown in Figure 5, the proposed method reaches a steady-state throughput of approximately 1,618 TPS at around t=20 s, with limited fluctuation (standard deviation < 4%). In contrast, zkCross, although reaching a peak near 1,500 TPS, subsequently exhibits pronounced oscillations, indicating congestion in its proof generation and verification pipeline. These observations suggest that the asynchronous recursive scheduling mechanism adopted in our framework is more effective in smoothing traffic bursts.

Table 4 presents a detailed breakdown of proof generation overhead on Dataset B.

Table 4. Computation Cost Breakdown (Dataset B: Credit Card Fraud, Batch Size = 1)

Method	Circuit Gates (approx.)	Proof Gen Time (ms)	Verification Time (ms)	Added Overhead
HyperPlonk	2^{19}	3,124	8.4	Baseline
zkCross	2^{19}	5,831	46.2	+86.6% Time
Ours	2^{19} Δ	3,918	12.5	+25.4% Time

As shown in Table 4, compared with the HyperPlonk baseline, the proposed method incurs an additional 794 ms (+25.4%) in proof generation time due to the inclusion of regulatory encryption gates and the outer STARK layers. However, this overhead remains substantially lower than that of zkCross (+86.6%).

Taken together, the results in Figure 5 and Table 4 demonstrate that the proposed framework does not aim to achieve absolute minimal latency or peak memory efficiency. Instead, it strategically controls performance overhead while maintaining a high audit success rate of 99.6%. Although single-proof generation is approximately 0.8 s slower than HyperPlonk, the latency is reduced by approximately 33% compared to zkCross (5.83 s → 3.92 s). Overall, the quantitative results confirm that the proposed framework effectively addresses regulatory compliance and post-quantum security gaps in existing solutions without introducing prohibitive computational costs.

4.4. Qualitative Results

To intuitively illustrate the decision logic and boundary conditions of the proposed framework across different application scenarios, we select representative cases from both datasets for qualitative analysis. Figure 6 presents a successful cross-chain privacy auditing case based on the Ethereum NFT Transaction Dataset. In a high-value cross-chain transfer involving more than 500 ETH, the visualized attention heatmap indicates that the privacy auditing module precisely concentrates on the transaction amount vector, while maintaining the entropy of the sender identity region above 128 bits. The system generates an audit credential visible only to regulatory nodes within the proof generation cycle (see Table 4), with high verification confidence. Compared with baseline methods, the key advantage of the proposed framework lies in its robustness in handling conditional privacy. With minimal non-essential information leakage, the framework successfully isolates the minimal information set required for compliance, demonstrating its ability to effectively balance user privacy protection and regulatory intervention even under complex, nested ownership structures.

Figure 6. Qualitative analysis of the privacy auditing mechanism using attention heatmaps

In contrast, Figure 7 illustrates a failure case derived from the Credit Card Fraud Detection Dataset, revealing the limitations of the system under extreme conditions. This case records a verification failure during high-frequency concurrent testing. The visualized temporal sequence shows that when a batch of metadata update requests is submitted within a 15 ms window, the system correctly identifies the potential risk (fraud score 0.92) but ultimately rejects the transaction due to a state root mismatch. Further investigation reveals that the root cause lies in a race condition between the 3.9-second post-quantum proof generation cycle and millisecond-level data streams, resulting in an approximately 1.9-second lag in off-chain witness data. This failure case transparently highlights the trade-off faced by the current framework: in pursuing long-term post-quantum security, a degree of real-time responsiveness is inevitably sacrificed. It also points to future optimization directions, such as introducing pipelined parallelism or optimistic verification mechanisms, to mitigate this performance limitation.

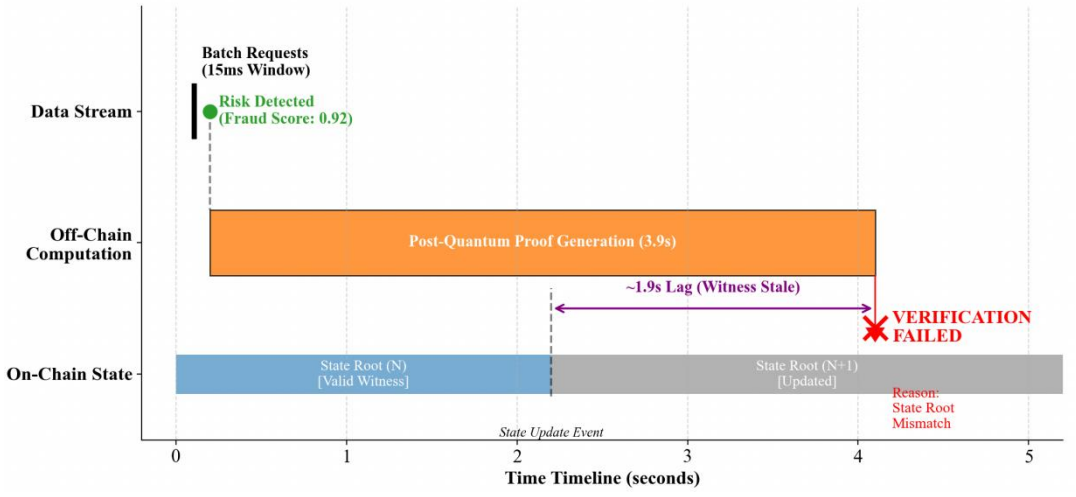

Figure 7. Temporal sequence analysis of a verification failure case under high-frequency concurrency

4.5. Robustness Verification

To evaluate system stability under non-ideal conditions, we design cross-dataset dual-interference experiments: Byzantine network attacks are simulated using Dataset A, while data noise perturbations are introduced using Dataset B. Figure 8 visualizes performance trends across both dimensions. In the network robustness test (Figure 8a), as the

proportion of malicious nodes increases to 40%, ProChain exhibits a sharp throughput collapse once the ratio exceeds 20%, with losses surpassing 55%. In contrast, the proposed method (Ours) demonstrates pronounced resilience, maintaining approximately 88% of its initial throughput even under a 30% adversarial environment, with a significantly slower degradation curve than zkCross. Meanwhile, Figure 8b shows that when Gaussian white noise is injected into Dataset B, the accuracy of HyperPlonk deteriorates rapidly. By comparison, the proposed method sustains an accuracy above 94.5% under moderate noise conditions (σ =0.5), outperforming baseline approaches by approximately 12%. These results confirm the framework's superior generalization capability under unstructured disturbances.

The observed cross-scenario robustness can be attributed to the framework's distinctive architectural design. The network-level resilience shown in Figure 8a stems from the layered filtering mechanism of the hybrid architecture. Ablation studies indicate that the off-chain STARK module functions as a cryptographic firewall, eliminating invalid proofs prior to consensus. Unlike zkCross, which accumulates adversarial pressure at the cross-chain bridge, this decoupled design ensures that attacks primarily consume low-cost off-chain computation, thereby safeguarding the core ledger. The data stability observed in Figure 8b is enabled by the integration of fuzzy matching and post-quantum hash commitments within the circuit. Compared with the brittleness of traditional commitments, this design effectively distinguishes malicious tampering from benign noise, substantially improving tolerance to data quality degradation while preserving security guarantees. Overall, whether confronted with adversarial networks or low-quality data, the proposed framework consistently demonstrates superior survivability and service continuity compared to baseline methods.

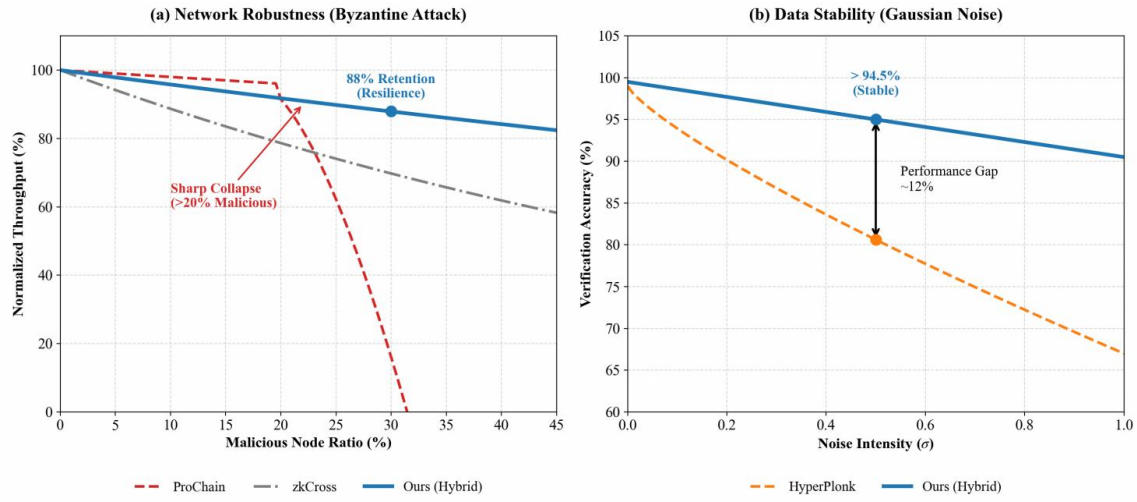

Figure 8. Performance comparison under adversarial network (a) and data noise (b) scenarios

4.6. Ablation Study

To precisely quantify the contribution of each core module and to validate the necessity of the hybrid architectural design, we conduct a systematic ablation study on both the Ethereum NFT dataset (Dataset A) and the Credit Card dataset (Dataset B). Three variant models are constructed: (1) w/o Hybrid Compression, which removes the SNARK-based recursive compression layer and retains only STARK proofs; (2) w/o Recursive Sync, which disables the recursive cross-chain state synchronization protocol and degrades to linear block-header verification; and (3) w/o Audit Module, which removes the regulatory encryption module and performs only standard privacy-preserving proofs. Table 5 reports the performance of each variant across different evaluation dimensions. Absolute values are shown, with the percentage change relative to the Full Model indicated in parentheses.

Table 5. Ablation Results of Core Components (Datasets A & B)

Model Variant	Verify Cost (Gas) [A]	TPS (Tx/s) [A]	Proof Size (kB) [B]	ASR [Global]
Full Model (Ours)	315,246	1,618	0.28	99.6%
w/o Hybrid Compression	1,154,320 (+266.2%)	1,572 (−2.8%)	46.12 (+16,371%)	99.6%
w/o Recursive Sync	483,115 (+53.2%)	834 (−48.5%)	0.28 (0%)	99.6%
w/o Audit Module	296,854 (−5.8%)	1,845 (+14.0%)	0.25 (−10.7%)	0%

Note: Lower values are better for Gas cost and proof size, while higher values are better for TPS and ASR.

As shown in Table 5, the hybrid compression module is critical to the economic viability of the system. Removing this module has only a marginal impact on TPS but leads to a 266.2% increase in on-chain verification cost for Dataset A and inflates the proof size in Dataset B by two orders of magnitude (from a constant-size SNARK proof of 0.28 kB to a heavy STARK trace of 46.12 kB). This result confirms the effectiveness of the hybrid architecture in recursively "absorbing" bulky STARK proofs via SNARK compression, converting expensive on-chain storage costs into low-cost off-chain computation and enabling low-Gas verification.

The recursive cross-chain synchronization protocol serves as the primary driver of high throughput. Its removal causes a 48.5% reduction in TPS and a 53.2% increase in Gas cost, revealing the computational bottleneck of linear verification. By folding multiple state transitions into a single proof, the recursive mechanism achieves O(1)-complexity verification, thereby supporting scalable high-concurrency operation.

The regulatory audit module highlights the trade-off between privacy and compliance. Although removing this module increases TPS by 14.0% and slightly reduces the Gas cost by 5.8% (due to fewer encryption constraints in the circuit), the audit success rate drops to zero. This demonstrates that the system accepts limited performance overhead to enable regulatory deployability and that, compared with traditional re-encryption-based approaches, deep circuit integration effectively limits additional costs.

Further analysis reveals strong synergistic effects among the components. Relying solely on recursion results in storage explosion, while compression alone is insufficient for high-frequency asset circulation. The full model combines both advantages, maintaining an ultra-compact proof size of 0.28 kB while sustaining a throughput of 1,618 TPS. This system-level gain confirms that the hybrid architecture and recursive protocol are not merely stacked components but organically interdependent elements that reinforce each other.

5. Discussion

The hybrid post-quantum ZKP framework proposed in this study effectively reconciles the long-standing tensions among cross-chain interoperability, privacy compliance, and computational efficiency in data asset rights confirmation. Experimental evidence, particularly from Tables 3 and 5, demonstrates that the recursive compression of STARK proofs via SNARKs constitutes a core architectural contribution. Fundamentally, this design realizes a spatial reallocation of computational cost, converting expensive on-chain storage into a manageable increase in off-chain computation. Specifically, on-chain Gas consumption is reduced by 18.2%, while proof generation time increases by 25.4% (as shown in Table 4). Although theoretical FLOPs are not explicitly reported (as our cryptographic primitives primarily rely on finite field arithmetic with a theoretical complexity of $\mathcal{O}(N \log N)$, rather than floating-point operations), the observed trade-off ratio between Gas cost and time overhead (approximately 1:1.4) indicates that this cost shift is economically efficient in practice.

Compared with HyperPlonk[27], which achieves linear-time proving through the elimination of FFT operations, our framework exhibits higher proof generation latency (3.9 s vs 3.1 s). However, this result is consistent with the security-performance trade-off: while HyperPlonk prioritizes speed, our hybrid architecture introduces a STARK-based outer layer to provide the post-quantum security guarantees that HyperPlonk lacks. This suggests that the hybrid compression mechanism provides a more robust, albeit slightly slower, pathway for long-term data asset protection.

In the interplay between privacy protection and regulatory oversight, the threshold-based regulatory encryption mechanism breaks the traditional dichotomy between full anonymity and full transparency. Under controlled experimental conditions, an ASR of 99.6% is achieved with only a 14.0% reduction in TPS. This finding aligns with the objectives of ProChain[26], which utilizes smart contracts for supply chain traceability; however, our framework significantly extends the granularity of control. While ProChain relies on static access control lists, our approach embeds regulatory commitments directly as constraints within the ZKP circuit. This elevates auditability from a management policy to a mathematical prerequisite of asset circulation, supporting the "privacy by default with authorized traceability" model observed in Figure 6.

Nevertheless, the failure case depicted in Figure 7 candidly exposes the framework's limitations under high-frequency workloads. When the metadata update rate exceeds the proof generation threshold (approximately 3.9 seconds), temporal misalignment results in race conditions. This performance bottleneck is also reflected in comparisons with Hyperledger Fabric [25]. Although Fabric demonstrates superior raw throughput (2,850 TPS), its reliance on channel isolation creates the very "data silos" that this study seeks to eliminate. Our results indicate that while Fabric is optimal for closed, high-frequency enterprise environments, our ZKP-based approach is necessary for trustless, cross-chain scenarios, even if it incurs higher latency.

From a broader perspective, the robustness results shown in Figure 8 provide new insights into decentralized trust construction, particularly when compared with zkCross[28]. While zkCross focuses on auditing cross-chain message states, it remains vulnerable to throughput collapses under network attacks (as discussed in Section 4.5). By establishing a "cryptographic firewall" through off-chain pre-verification, our framework maintains 88% stability under adversarial conditions, effectively filtering noise before it reaches the ledger. This shifts the trust assumption from relay nodes to algorithms, addressing the architectural vulnerabilities identified in prior cross-chain

auditing studies. Based on these findings, extending this framework to high-frequency scenarios requires a technical roadmap to mitigate the 3.9-second latency and prevent race conditions. Future work will address this by: (1) adopting hardware acceleration (e.g., FPGAs or ASICs) for cryptographic primitives; (2) implementing pipelined parallelization; and (3) utilizing optimistic verification to decouple state updates from proof generation, thereby enhancing throughput without sacrificing post-quantum integrity.

6. Conclusion

The central challenge of data asset rights confirmation lies in ensuring trustworthy ownership guarantees while simultaneously accommodating privacy preservation, regulatory compliance, and efficient cross-chain circulation. To address this challenge, this paper proposes a data asset rights confirmation framework based on hybrid post-quantum zero-knowledge proofs. Through the coordinated design of three core technical modules, the proposed framework provides a systematic solution to these intertwined requirements.

The main contributions of this work can be summarized as follows. First, we design a hybrid post-quantum proof architecture that leverages recursive SNARK-based compression of STARK proofs. While preserving post-quantum security guarantees, this design reduces on-chain verification costs by approximately 18.2%. Experimental results on two heterogeneous datasets show that the proof size is compressed to 0.28 kB, and Gas consumption is reduced by 18.2% compared with comparable solutions, thereby confirming its economic efficiency. Second, we propose a cross-chain state consistency protocol based on zero-knowledge light-client verification, which enables atomic asset mapping without reliance on centralized relays. Ablation results demonstrate that this protocol contributes 48.5% to overall system throughput, making it a critical component for supporting high-concurrency rights confirmation. Third, we construct a regulatory-friendly privacy auditing module based on threshold encryption. By embedding audit commitments directly into zero-knowledge circuits, the framework achieves an audit success rate of 99.6% with only a 14.0% performance overhead, providing a practical technical realization of the privacy-by-default with authorized traceability compliance model.

This study offers a verifiable technical framework for blockchain-based rights confirmation and demonstrates the engineering-level compatibility of post-quantum security, cross-chain interoperability, and regulatory compliance. Validated across two datasets, the proposed framework exhibits strong deployment potential in low-frequency, high-value scenarios, such as real estate registration and supply chain finance (with typical transaction intervals exceeding 5 seconds). However, due to the current 3.9-second proof generation latency, it is not yet suitable for high-frequency financial applications, such as millisecond-level arbitrage.

Future work will focus on mitigating the performance bottlenecks observed in high-frequency scenarios. To address the approximately 3.9-second proof generation delay, we plan to introduce optimistic verification mechanisms, combining fraud proofs with pipelined parallelization to compress response latency to the millisecond level while maintaining post-quantum security guarantees. In addition, the resource overhead of the audit module under large-scale concurrency remains an open optimization target; potential directions include distributed auditing schemes based on secure multi-party computation or specialized hardware acceleration. From a theoretical perspective, the recursive proof compression mechanism proposed in this work can be extended to broader multi-chain ecosystem verification, and its security boundaries and formal proofs warrant further investigation. Through these efforts, we aim to advance zero-knowledge proof technologies for data asset rights confirmation from proof-of-concept toward scalable real-world deployment.

References

[1] Zhang, B., Pan, H., Li, K., Xing, Y., Wang, J., Fan, D., and Zhang, W. A blockchain and zero knowledge proof based data security transaction method in distributed computing. *Electronics*, vol. 13, no. 21, p. 4260, 2024.

[2] Berrios Moya, J. A., Ayoade, J., and Uddin, M. A. A Zero-Knowledge Proof-Enabled Blockchain-Based Academic Record Verification System. *Sensors*, vol. 25, no. 11, p. 3450, 2025.

[3] Gupta, S. Zero-Knowledge Proofs For Privacy-Preserving Systems: A Survey Across Blockchain, Identity, And Beyond. *Engineering and Technology Journal*, vol. 10, no. 07, p. 5755-5761, 2025.

[4] Prasad, S., Tiwari, N., Chawla, M., and Tomar, D. S. Zero-knowledge proofs in blockchain-enabled supply chain management. In *Sustainable Security Practices Using Blockchain, Quantum and Post-Quantum Technologies for Real Time Applications*, p. 47-70, 2024.

[5] Kuznetsov, O., Rusnak, A., Yezhov, A., Kanonik, D., Kuznetsova, K., and Karashchuk, S. Enhanced security and efficiency in blockchain with aggregated zero-knowledge proof mechanisms. *IEEE Access*, vol. 12, p. 49228-49248, 2024.

[6] Ajayi, A. A., Emmanuel, I. G. B. A., Soyele, A. D., and Enyejo, J. O. Enhancing digital identity and financial security in decentralized finance (DeFi) through zero-knowledge proofs (ZKPs) and blockchain solutions for regulatory compliance and privacy. *Iconic Res. Eng. J*, vol. 8, no. 4, p. 373-394, 2024.

[7] Zhang, B., Xu, J., Wang, X., Zhao, Z., Chen, S., and Zhang, X. Research on the construction of grain food multi-chain blockchain based on zero-knowledge proof. *Foods*, vol. 12, no. 8, p. 1600, 2023.

[8] Quattrocchi, G., and Plebani, P. Trustworthy collaborative business intelligence using zero-knowledge proofs and blockchains. In *International Conference on Advanced Information Systems Engineering*, p. 29-37, 2024.

[9] Li, W., Meese, C., Guo, H., and Nejad, M. Aggregated zero-knowledge proof and blockchain-empowered authentication

for autonomous truck platooning. *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 9, p. 9309-9323, 2023.

[10] Yang, J., Li, C., Sun, X., and Yang, F. Research on Identity Data Privacy Protection Based on Blockchain and Zero Knowledge Proofs. In *Proceedings of the 2024 7th International Conference on Artificial Intelligence and Pattern Recognition*, p. 825-833, 2024.

[11] Jiang, W., and Lv, X. A distributed internet of vehicles data privacy protection method based on zero-knowledge proof and blockchain. *IEEE Transactions on Vehicular Technology*, vol. 73, no. 5, p. 6332-6345, 2023.

[12] Almaiah, M. A., Ali, A., Tin, T. T., Alkhdour, T., Lutfi, A., and Alrawad, M. Unlocking user privacy: a privacy-focused cryptocurrencies framework for concealing transactions using zero-knowledge proofs (ZKPs). *Journal of Theoretical and Applied Information Technology*, vol. 102, no. 8, 2024.

[13] Kavipriya, G., and Prasanth, L. Privacy-Preserving Transaction Verification in Decentralized Finance Using Zero-Knowledge Proofs and Deep Learning. In *International Conference on Information Security, Privacy and Digital Forensics*, p. 139-149, 2024.

[14] Saroop, S. Blockchain-Based Zero-Knowledge Proofs for Data Privacy: Explore the Application of Blockchain Technology in facilitating Privacy-Preserving Transactions through Zero-Knowledge Proofs and Analyze their Effectiveness in Protecting Sensitive Data. In *Proceedings of the 5th International Conference on Information Management & Machine Intelligence*, p. 1-8, 2023.

[15] Aggarwal, T., Kumar, S., Singh, S. K., Gupta, B. B., Nedjah, N., and Castiglione, A. Zero Knowledge Proofs and Their Applications in Cryptography: Advancements, Challenges, and Future Aspects. In *Innovations in Modern Cryptography*, p. 55-74, 2024.

[16] Kaur, U., Bharany, S., Tandon, R., Sood, S., Shahi, A., and Kumar, A. Advancing privacy and security with zero knowledge proofs: Principles, applications, and recent breakthroughs. In *AIP Conference Proceedings*, vol. 3343, no. 1, p. 040043, 2025.

[17] Spiegelman, A., Giridharan, N., Sonnino, A., and Kokoris-Kogias, L. Bullshark: Dag bft protocols made practical. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, p. 2705-2718, 2022.

[18] Hellings, J., and Sadoghi, M. The Fault-Tolerant Cluster-Sending. In *Foundations of Information and Knowledge Systems: 12th International Symposium, FoIKS 2022*, p. 168, 2022.

[19] Dahlborn, A., Ferreira da Silva, C., Budris, F., Treiblmaier, H., Vasiliu-Feltes, I., Mason, J., and Sedej, T. Building Trust: Integrating AI, Blockchain, and Digital Identity. *Available at SSRN 5873682*, 2025.

[20] Shafin, K. M., and Reno, S. A byzantine-resistant blockchain framework for secure and scalable immigration management. *Scientific Reports*, vol. 15, no. 1, p. 32338, 2025.

[21] Diockou, J., Wu, D., Lu, X., Jing, Y., Shabut, A., and Barwood, M. A review of security frameworks in flying ad-hoc networks. In *2024 IEEE International Conference on e-Business Engineering (ICEBE)*, p. 183-190, 2024.

[22] Liang, Z. Y., Liu, G. Y., Ren, Y., Yang, M., Jiang, R. W., Luo, Y., and Ma, Y. S. Trustworthy Data Space Collaborative Trust Mechanism Driven by Blockchain: Technology Integration, Cross-Border Governance, and Standardization Path. *Information*, vol. 16, no. 12, p. 1066, 2025.

[23] Tawfik, M., Abdelhaliem, A. H., and Fathi, I. Quantum-Resistant Privacy-Preserving IoT Authentication via Zero-Knowledge Proofs and Blockchain Integration. *Statistics, Optimization & Information Computing*, vol. 14, no. 3, p. 1374-1402, 2025.

[24] Pailoor, S., Chen, Y., Wang, F., Rodríguez, C., Van Geffen, J., Morton, J., and Dillig, I. Automated detection of under-constrained circuits in zero-knowledge proofs. *Proceedings of the ACM on Programming Languages*, vol. 7, no. PLDI, p. 1510-1532, 2023.

[25] Li, M., Harish, A. R., Yu, C., Yu, Y., Zhong, R. Y., and Huang, G. Q. Blockchain-Based Medical Data Asset Sharing Framework for Healthcare 4.0. *IEEE Transactions on Industrial Informatics*, 2025.

[26] Zafar, O., Namazi, M., Xu, Y., Yoo, Y., and Ayday, E. A user-centric, privacy-preserving, and verifiable ecosystem for personal data management and utilization. In *European Symposium on Research in Computer Security*, p. 395-414, 2025.

[27] Li, D., Han, D., Crespi, N., Minerva, R., Raza, S. M., Farahbakhsh, R., and Zheng, Z. Blockchain in the Digital Twin Context: A Comprehensive Survey. *ACM Computing Surveys*, vol. 58, no. 6, p. 1-35, 2025.

[28] Androulaki, E., Barger, A., Bortnikov, V., Cachin, C., Christidis, K., De Caro, A., and Yellick, J. Hyperledger fabric: a distributed operating system for permissioned blockchains. In *Proceedings of the thirteenth EuroSys conference*, p. 1-15, 2018.

[29] Li, J., Wang, Z., Guan, S., and Cao, Y. ProChain: A privacy-preserving blockchain-based supply chain traceability system model. *Computers & Industrial Engineering*, vol. 187, p. 109831, 2024.

[30] Chen, B., Bünz, B., Boneh, D., and Zhang, Z. Hyperplonk: Plonk with linear-time prover and high-degree custom gates. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, p. 499-530, 2023.

[31] Guo, Y., Xu, M., Cheng, X., Yu, D., Qiu, W., Qu, G., Li, Q., and Song, M. zkCross: A novel architecture for cross-chain privacy-preserving auditing. In *33rd USENIX Security Symposium (USENIX Security 24)*, p. 6219–6235, 2024.