

A Privacy-Preserving Data Sharing Framework Driven by Blockchain for Smart Childcare Subsidy Systems

Ruiyang Li¹, Xiaomei Liu^{1*}

¹School of Public Administration, Dongbei University of Finance and Economics, Dalian 116025, China

Abstract

With the continuing development of digital government and public childcare-assistance programs, cross-departmental subsidy services must jointly address sensitive-data minimization, dynamic authorization, fraudulent-application identification, and accountable auditing. This paper proposes a blockchain-driven privacy-preserving data-sharing framework that combines a consortium blockchain, on-chain evidence with off-chain ciphertext storage, zero-knowledge eligibility proofs, attribute-based encryption, proxy re-encryption, secure aggregation, federated learning, and smart-contract auditing. The contribution is a scenario-specific system integration and governance design rather than a new cryptographic primitive. In addition to the original Hyperledger Fabric feasibility test, the revision adds literature-grounded validation against peer-reviewed studies of childcare-subsidy administration and administrative-data evaluation, component implementation-status disclosure, a 2048-bit proof-of-knowledge precomputation microbenchmark, non-IID federated-learning tests, and colluding-node poisoning tests. The original prototype reaches 718 TPS and 236 ms average latency at 400 concurrent requests, 91.8% federated-learning accuracy, and 138 ms on-chain index retrieval for 20,000 audit records. In the added five-run tests, strong non-IID data retain an F1 score of 82.69% $\pm$ 0.09%; adaptive clipping and robust aggregation retain F1 scores of 84.04% and 84.14% under one and two colluding sign-flip clients, respectively. Offline commitment precomputation reduces the online arithmetic time for four eligibility subproofs from 8.03 $\pm$ 0.34 ms to 0.06 $\pm$ 0.01 ms. These results support the feasibility of the integrated framework while clarifying that production deployment still requires authorized local data, audited cryptographic libraries, and multi-orderer field validation.

Keywords: Smart civil affairs, childcare subsidy, blockchain, privacy protection, zero-knowledge proof, attribute-based encryption, smart contract, audit traceability

Received on 16 June 2026, accepted on 17 August 2026, published on 07 October 2026

Copyright © 2026 Ruiyang Li *et al.*, licensed to EAI. This is an open access article distributed under the terms of the [CC BY-NC-SA 4.0](#), which permits copying, redistributing, remixing, transformation, and building upon the material in any medium so long as the original work is properly cited.

doi: 10.4108/eetsis.13557

1. Introduction

1.1 Research Background and Significance

Peer-reviewed studies of state-level childcare subsidy and assistance programs describe regulated public-service workflows involving application intake, eligibility determination and redetermination, participation continuity, subsidy administration, and program monitoring [1-3]. These

studies also use administrative records covering family and child characteristics, eligibility and participation, provider information, payments, and program performance. In this manuscript, the term 'lead subsidy authority' denotes a configurable administrative role in the prototype. The role can be mapped to the competent agency under different national or local institutional arrangements, preventing the technical architecture from being tied to a single jurisdiction.

However, existing childcare subsidy systems still mainly rely on centralized databases, departmental interface calls,

*Corresponding author. Email: ltthy@163.com

and manual material review. They suffer from inconsistent data standards, complex authorization processes, unclear sharing responsibility boundaries, and insufficient review efficiency. Meanwhile, childcare subsidy services involve highly sensitive information such as minors' identities, family relationships, bank accounts, and subsidy receipt records. If plaintext transmission, centralized storage, or coarse-grained authorization is adopted, risks such as unauthorized queries, interface abuse, database leakage, and third-party attacks may arise.

The integration of blockchain and privacy computing provides a feasible solution to these problems. Blockchain can provide tamper-proof evidence storage, timestamp recording, automatic smart contract execution, and full-process audit capabilities for multi-department collaboration. However, raw sensitive data should not be directly uploaded to the blockchain. Therefore, this paper adopts a collaborative model of "on-chain evidence storage, off-chain ciphertext storage, on-chain authorization audit, and off-chain privacy computing", so that only hash digests, permission credentials, and audit records are stored on-chain, while raw data remain in local departmental systems or secure off-chain storage systems.

Privacy computing further enhances the secure sharing capability of blockchain systems. Zero-knowledge proof can verify subsidy eligibility without disclosing complete household registration, birth certificate, or guardianship relationship data; homomorphic encryption and secure aggregation can support ciphertext statistics and risk scoring; attribute-based encryption can implement fine-grained access control based on departmental responsibilities, job roles, business stages, and authorization periods; and federated learning can identify duplicate applications and abnormal applications without moving data out of local domains.

Based on this, this paper focuses on data silos, privacy leakage, fraudulent applications, and untraceable processes in smart civil affairs childcare subsidy systems, and constructs a blockchain-driven privacy-preserving data sharing framework. The framework realizes a subsidy data governance model in which "data are usable but invisible, eligibility is provable without disclosure, permissions are controllable and revocable, and processes are trusted and traceable", thereby providing technical support for secure collaboration, accurate review, and trusted supervision in smart civil affairs systems.

1.2 Research Status at Home and Abroad

1.2.1 Research Status of Smart Government Affairs and Livelihood Subsidy Systems

Digital-government research has moved from online service delivery toward interdepartmental data exchange, automated eligibility review, program-integrity control, and accountable governance. Peer-reviewed studies of state-level childcare subsidy and assistance programs show that eligibility and recertification rules, reporting burdens, participation stability, and administrative-data quality are central to program operation and evaluation [1-3]. Collectively, these studies

motivate a data model covering family, child, eligibility, participation, provider, and program-performance information.

1.2.2 Research on Blockchain Applications in Government Data Sharing

General blockchain surveys establish the technical characteristics, application landscape, and open challenges of distributed ledgers [4-6]. Public-sector and broader organizational blockchain studies examine governance, institutional adoption, and accountability-related factors [7,10-12], while secure e-government and government-data-sharing studies address privacy protection, authorization, trusted evidence, and interdepartmental sharing [8,13,14]. The Service-On-Chain approach keeps sensitive data off chain while formalizing departmental service agreements through smart contracts [13], whereas representative secure e-government frameworks combine permissioned blockchain with privacy and security controls [8,14]. Smart-contract access-control research demonstrates how static and dynamic authorization policies can be implemented through blockchain-based contracts [9]. Broader IoT-security research highlights the enlarged attack surface of connected infrastructures [15], while privacy-sensitive healthcare blockchain studies emphasize privacy-preserving handling of sensitive records and limiting unnecessary disclosure [16,17].

1.2.3 Research Progress of Federated Learning and Privacy Computing in Trusted Data Collaboration

Federated learning supports collaborative modeling without centralizing raw records [18-22]. Blockchain-enabled federated-learning surveys emphasize ledger-assisted auditability, model-update integrity, and governance [23,24], while broader federated-learning security reviews identify poisoning, inference, privacy leakage, and robustness as unresolved issues [25,26]. Secure aggregation protects individual updates [27], and recent scalable-information-systems research additionally combines input verification and dropout resistance to strengthen aggregation against invalid or unavailable participants [28]. The reported experiments in this paper use horizontal federated logistic regression across regional subsidy-service nodes with a common risk-feature schema; alternative federated model classes and configurations are documented in the literature [29,30] and may be considered in future work.

1.2.4 Research Progress of Attribute-Based Encryption, Zero-Knowledge Proof, and Dynamic Access Control

Fine-grained government-data authorization requires more than static roles. Attribute-based encryption [32,33], building on earlier attribute-oriented identity encryption such as fuzzy identity-based encryption [31], is complemented by proxy re-encryption [34], zero-knowledge proofs [35-37], homomorphic encryption [38,39], group signatures [40], encrypted-data search [41], differential privacy [42,43], and multi-layer access-control models for cross-organizational sensitive-data sharing [44]. These mechanisms have different security assumptions, lifecycle rules, and computational

costs; consequently, an integrated system must explicitly state which components are implemented, which are interface-level options, and which metrics are actually measured.

Among the representative studies reviewed above [7-14,23,24,44], few simultaneously address minimum-disclosure eligibility verification, dynamic authorization, cross-department risk identification, fund-disbursement

tracking, correction and revocation, and full-process auditing within a childcare-subsidy workflow. The novelty of this study lies in systematic integration and scenario-specific adaptation of established technologies, not in claiming a new cryptographic primitive. Table 1 positions the proposed framework against representative e-government and blockchain-federated-learning research.

Table 1. Positioning against representative government-data-sharing frameworks

Study	Application scenario	Privacy and access control	Fraud / joint risk detection	Auditability
Ølmes et al. [7]	General government information sharing	Conceptual/editorial governance discussion	Not included	Blockchain implications
Piao et al. [13]	Interdepartmental government data sharing	Off-chain data; smart-contract service agreements	Not included	Service-state evidence
Elisa et al. [8,14]	General secure e-government	Permissioned blockchain and privacy controls	Limited / not subsidy-specific	Transaction traceability
Blockchain-FL studies [23,24]	Cross-domain collaborative learning	Model-update privacy and ledger audit	Generic model risk	Training-process audit
Proposed framework	Childcare subsidy application, review, payment, and supervision	ZKP + ABE + PRE + dynamic authorization + off-chain ciphertext	Federated duplicate/fraud risk identification	Application-to-disbursement full-process audit
Zhang et al. [9]	IoT smart-contract access control	Smart-contract policy validation and dynamic authorization	Misbehavior judgment; not subsidy-specific fraud	On-chain contract state and access-control events
Chenthara et al. [44]	Cross-organizational electronic-health sharing	Multi-layer / attribute-aware access control	Not included	Provenance-supported data integrity

1.3 Research Content and Contributions

To address data silos, privacy leakage, manual-review burden, fraudulent claims, and weak process traceability, this paper proposes a scenario-adapted integration framework. The work does not introduce a new encryption, signature, proof, or federated-learning primitive. Its contribution is the co-design of business roles, data lifecycle, privacy modules, smart contracts, risk-control workflow, and audit evidence for a regulated childcare-subsidy service.

The revised study follows the line of policy and business-process mapping, system architecture, privacy and security mechanism design, implementation-scope disclosure, and reproducible evaluation. The main contributions are:

- An on-chain/off-chain collaborative data sharing architecture for smart civil affairs childcare subsidy systems is proposed to enable trusted evidence storage, status synchronization, and full-process auditing of cross-departmental subsidy data;
- A layered privacy protection mechanism integrating zero-knowledge proof, attribute-based encryption, proxy re-encryption, certificate-based pseudonymous identity management, and an optional group-signature interface

is designed to achieve minimum-disclosure eligibility verification and fine-grained dynamic authorization;

- A federated learning-based fraudulent application identification and smart contract collaboration model with blockchain auditing is constructed to enable cross-departmental joint risk control and accountable supervision without moving raw data out of local domains.

2. Related Theories and Technical Foundations

2.1 Childcare Subsidy Business Scenario and Security Requirement Modeling

A childcare-subsidy system is a regulated, multi-entity public-service workflow involving applicant families, a lead subsidy authority, service providers or supporting verification units, payment institutions, and oversight bodies. Prior studies of state-level childcare subsidy and assistance programs identify application and eligibility management, periodic redetermination, participation monitoring, provider records, and program evaluation as core operational elements

[1-3]. The prototype abstracts these roles so that the workflow owner and supporting entities can be configured for different institutional arrangements.

For technical modeling, the process is condensed into six stages: application, evidence storage, cross-department verification, eligibility and risk review, fund disbursement, and post-event audit. Each stage invokes only the minimum

necessary data, verifies authorization, updates status, and produces an auditable event.

Figure 1 summarizes the principal business, verification, and audit paths. Solid lines indicate the main business flow, dashed lines indicate cross-department verification, and dash-dotted lines indicate supervision. The diagram is an adaptable technical abstraction rather than a claim that every locality assigns identical duties to the same department.




Figure 1. Flowchart and Data Sharing Relationship Diagram of Smart Civil Affairs Childcare Subsidy Business

To describe the multi-entity collaborative relationship in the childcare subsidy system, the set of participants is defined as:

$$P = \{P_f, P_m, P_p, P_h, P_s, P_i, P_b, P_a\} \quad (1)$$

where P_f denotes the applicant family, P_m denotes the lead subsidy authority, P_p denotes the public security department, P_h denotes the health commission, P_s denotes the human resources and social security/medical insurance department, P_i denotes the finance department, P_b denotes the financial institution, and P_a denotes the regulatory audit department. Different entities play different roles in the system and are subject to both data sharing relationships and permission constraints as well as regulatory responsibility boundaries.

The data set involved in the childcare subsidy system can be expressed as:

$$D = \{D_{id}, D_{birth}, D_{house}, D_{guard}, D_{social}, D_{subsidy}, D_{bank}, D_{audit}\} \quad (2)$$

where D_{id} denotes identity data, D_{birth} denotes birth medical certificate data, D_{house} denotes household registration data, D_{guard} denotes guardianship relationship data, D_{social} denotes social security/medical insurance data, $D_{subsidy}$ denotes subsidy receipt records, D_{bank} denotes bank account and fund disbursement data, and D_{audit} denotes audit log data. These data have different sensitivity levels and access conditions during business review, so they cannot be managed by a single plaintext sharing or static authorization method.

From the perspective of security requirements, the smart civil affairs childcare subsidy system should satisfy the following goals. First, privacy protection: sensitive information such as minors' identities, family relationships, bank accounts, and subsidy records should not be obtained by unauthorized entities during cross-departmental circulation. Second, trusted sharing: data sources should be authentic, the sharing process should be verifiable, and review results should be non-repudiable. Third, traceability: a complete evidence chain should be formed for key stages such as

application submission, data invocation, permission granting, review confirmation, fund disbursement, and abnormal handling. Fourth, controllable permissions: fine-grained access control should be implemented according to departmental responsibilities, job roles, business stages, data categories, and authorization periods. Fifth, tamper resistance: review status, subsidy records, fund disbursement records, and audit logs should be protected from unilateral tampering.

Sixth, anti-fraud capability: risks such as duplicate applications, false certificates, cross-regional duplicate claims, identity impersonation, and abnormal fund flows should be identified.

Table 2 presents the mapping relationship between childcare subsidy business requirements, main risks, and key technologies.

Table 2. Matching Relationship Between Childcare Subsidy Business Requirements and Key Technologies

Business Process	Main Risks	Technical Requirements	Matched Technologies
Application and Material Submission	Identity leakage, material tampering, fraudulent submission	Identity protection, material integrity verification	Hash evidence; MSP/CA certificates; pseudonymous IDs; off-chain ciphertext; optional group-signature interface
Qualification Verification	Overexposure of birth, household registration and guardianship information	Minimum disclosure verification, cross-department trusted verification	Zero-knowledge proof, attribute-based encryption, proxy re-encryption
Multi-level Review	Inconsistent review standards, unauthorized access	Automatic process execution, fine-grained authorization	Smart contract, dynamic access control
Fund Disbursement	Inconsistent disbursement status, unclear fund flow	Status synchronization, traceable fund records	Blockchain deposit, disbursement contract
Violation Identification	Duplicate application, cross-region subsidy fraud	Cross-department joint risk control	Federated learning, secure aggregation
Post-event Supervision	Log tampering, unclear liability	Full-process audit, tamper-proof evidence chain	Distributed ledger, audit contract

As shown in Table 2, the security requirements of the childcare subsidy system cannot be fully addressed by a single technology. Instead, blockchain, privacy computing, access control, and smart contracts must work collaboratively.

2.2 Blockchain-Based Trusted Collaboration Mechanism

Blockchain is a distributed ledger jointly maintained by multiple authorized parties and can provide tamper-evident records, traceability, and smart-contract execution [4-7]. For childcare subsidy services, a consortium blockchain is used as an interdepartmental evidence and coordination layer rather than as a replacement for departmental business databases. Only hashes, indexes, signatures, authorization decisions, and audit events are placed on chain. Where encrypted indexes are required, searchable-encryption principles can be used to limit plaintext exposure [41].

$$R_i = \{\text{Hash}(C_i), \text{CID}_i, \text{Meta}_i, \text{Time}_i, \text{Sig}_i\} \quad (3)$$

where Hash (Ci) denotes the hash digest of the off-chain ciphertext data, CID_i denotes the data index generated by IPFS or an off-chain storage system, Meta_i denotes metadata such as data type, business stage, and sensitivity level, Time_i denotes the on-chain timestamp, and Sigi denotes the digital signature of the submitting entity. When a business node needs to access relevant data, the system first completes permission verification through a smart contract, then retrieves the off-chain ciphertext data according to CID_i, and finally completes authorized access through the corresponding key or re-encryption mechanism. This model balances data integrity, system scalability, and privacy protection, and is suitable for massive livelihood data scenarios such as childcare subsidies.

2.3 Key Privacy Protection Technologies

Privacy protection in the childcare subsidy system includes storage confidentiality, identity protection, eligibility verification, access control, ciphertext forwarding, collaborative modeling, and statistical release. The

framework therefore draws on established zero-knowledge proofs [35-37], homomorphic encryption [38,39], secure aggregation [27], attribute-based encryption [32,33], proxy re-encryption [34], group signatures [40], and optional differential-privacy mechanisms [42,43]. Fuzzy identity-based encryption [31] is cited as an early attribute-oriented precursor rather than as a CP-ABE scheme.

2.3.1 Federated Learning and Secure Aggregation

Federated learning retains applicant records locally at regional subsidy-service nodes and exchanges only model updates or intermediate results [18-22,30]. Secure aggregation prevents the coordinator from observing a single regional node's update [27,28]. In the reported prototype, every federated client uses the same risk-feature schema but holds a disjoint set of applications, making the experiment a horizontal federated-learning setting. Departmental source systems remain responsible for local verification and provide only authorized verification outputs to the regional risk node. Non-IID degradation, poisoning, inference, and client dropout remain relevant risks and are explicitly evaluated or discussed in the revision.

Suppose that the k -th regional subsidy-service node has a local dataset with the same feature schema as the other clients, its local model parameter is w_k , and there are K participating regional nodes in the system. The optimization objective of horizontal federated learning can be expressed as:

$$\min_w F(w) = \sum_{k=1}^K \frac{n_k}{n} F_k(w) \quad (4)$$

where $F_k(w)$ denotes the local loss function of the k -th regional node, n_k denotes the sample size of that node, and n denotes the total sample size. Each regional node uploads only model updates rather than applicant-level records. Secure aggregation is used so that the coordinator obtains only the aggregate and cannot inspect the raw update of an individual regional node.

In the proposed framework, blockchain can be used to record federated learning training rounds, model digests, node contributions, parameter verification results, and abnormal node behaviors, thereby addressing the insufficient trustworthiness of centralized aggregation nodes, the lack of training process auditability, and the difficulty of tracing malicious participants in traditional federated learning.

2.3.2 Attribute-Based Encryption and Dynamic Access Control

Fine-grained cryptographic authorization is supported by attribute-based encryption schemes [32,33], building on earlier identity-oriented constructions such as fuzzy identity-based encryption [31]. Complementary attribute-aware, multi-layer access-control models have also been investigated for cross-organizational sensitive-data sharing [44]. The access policy may combine organization, job responsibility, business stage, region, authorization period, and data sensitivity. Smart contracts record policy versions and authorization decisions, while the actual decryption

decision remains bound to the ciphertext policy and user attributes.

In the childcare subsidy system, visitor attributes may include department type, job responsibility, business stage, data category, administrative region, authorization period, and risk level. Let u denote the access subject, d denote the data to be accessed, t denote the business stage, $\text{Attr}(u)$ denote the subject attribute set, and $\text{Policy}(d,t)$ denote the access policy of data d at business stage t . Then the access control decision function can be defined as:

$$\text{Access}(u, d, t) = \begin{cases} 1, & \text{Attr}(u) \vdash \text{Policy}(d, t) \\ 0, & \text{otherwise} \end{cases} \quad (5)$$

where $\text{Access}(u,d,t)=1$ indicates that the access subject's attributes satisfy the access policy and the system permits access; $\text{Access}(u,d,t)=0$ indicates that the access request is rejected. Through this mechanism, front-line acceptance staff can only view material completeness and application status, lead-authority reviewers can view eligibility verification results, finance departments can view disbursement amounts and allocation status, and audit departments can view complete operation logs. This model can effectively reduce unauthorized access risks and support dynamic permission adjustment at different business stages.

2.3.3 Zero-Knowledge Proof, Homomorphic Encryption, and Privacy Verification

Zero-knowledge proofs allow a verifier to establish the validity of a statement without learning the underlying witness [36,37]. Privacy-preserving blockchain systems such as Zerocash illustrate an application of such techniques [35]. In this framework, the implemented prototype uses proof-of-knowledge and predicate-proof interfaces for age, guardianship, residence, and duplicate-claim conditions; it does not claim to invent a new proof system. Production deployment should use an audited library and standardized proof parameters.

Homomorphic encryption permits selected operations on ciphertext [38,39]. In the current manuscript it is retained as an optional interface for encrypted statistics and risk aggregation; it was not included in the end-to-end performance measurements. This distinction prevents conceptual components from being presented as fully benchmarked modules.

$$\text{Enc}(m_1) \oplus \text{Enc}(m_2) = \text{Enc}(m_1 + m_2) \quad (6)$$

In the childcare subsidy system, homomorphic encryption can be used for subsidy application quantity statistics, calculation of suspected duplicate claims, aggregation of cross-regional risk indicators, and verification of fiscal disbursement amounts. Compared with plaintext data aggregation, homomorphic encryption reduces data leakage risks during cross-departmental statistical processing. In practice, the system can select partially homomorphic encryption, additive homomorphic encryption, or lightweight schemes combined with secure aggregation according to

business complexity, so as to balance security and computational efficiency.

2.3.4 Group-Signature Interface and Proxy Re-Encryption

Group signatures provide signer anonymity and authorized opening under the security model of the selected scheme [40]. The cited short-group-signature construction does not by itself establish the dynamic joining and revocation properties needed for production deployment. Accordingly, the present prototype does not claim a fully dynamic group-signature implementation: it uses Fabric MSP/CA organizational certificates and pseudonymous application identifiers, while retaining only an interface for a future audited dynamic-group-signature module.

Proxy re-encryption enables ciphertext authorization conversion without revealing plaintext to the proxy [34]. The prototype evaluates the re-encryption service as a separate component and binds re-encryption-key issuance to a versioned authorization contract. Expired or revoked policies invalidate new access, while historical access events remain auditable.

2.4 Technical Applicability Analysis

The smart civil affairs childcare subsidy system has requirements for multi-entity collaboration, highly sensitive data sharing, strong regulation, and high traceability. A single technology cannot simultaneously satisfy trustworthiness, privacy, efficiency, and supervision requirements. Blockchain is suitable for trusted evidence storage of cross-departmental business records, automatic execution of review processes, and accountability tracing, but its on-chain transparency makes it unsuitable for directly storing raw sensitive data [4-7,13,14]. Federated learning is suitable for cross-departmental joint risk identification, but conventional federated learning can face centralized aggregation, model-

update privacy, poisoning, and auditability challenges [20,22-28]. Attribute-based encryption supports fine-grained authorization, while dynamic permission updates and ciphertext delegation may require complementary policy-management and proxy re-encryption mechanisms [31-34,44]. Zero-knowledge proofs support minimum-disclosure eligibility verification, although complex proofs may introduce computational overhead [35-37]. Homomorphic encryption supports selected computations on ciphertext, but computational cost remains an important deployment consideration [38,39]. Group signatures and proxy re-encryption can support signer anonymity with authorized opening and secure ciphertext forwarding, respectively [34,40]; production-grade dynamic membership or revocation would require an audited group-signature construction beyond the short-group-signature scheme cited in [40].

To further illustrate the applicability of blockchain and privacy-computing technologies to the childcare-subsidy process, Fig. 2 uses five macro-stages: application, eligibility verification, multi-level review, fund disbursement, and post-event supervision. This five-stage technology map is a higher-level abstraction of the six operational business stages defined in Section 2.1 and the seven system-processing steps in Section 3.5. In the application stage, the measured prototype uses hash evidence, Fabric MSP/CA authentication, pseudonymous application identifiers, and off-chain ciphertext storage. The eligibility-verification stage uses zero-knowledge proofs, attribute-based encryption, and proxy re-encryption. Smart contracts and dynamic access policies support review; blockchain records and disbursement contracts support fund-status synchronization; and horizontal federated learning, secure aggregation, and audit contracts support regional joint risk identification and accountability. Homomorphic encryption and group signatures are shown only as optional interfaces and are not included in the reported end-to-end benchmarks.




Figure 2. Adaptation of blockchain and privacy-preserving technologies to the childcare subsidy process

2.5 Implementation Scope and Component Status

To avoid conflating an integrated architecture with a complete production implementation, Table 3 distinguishes

measured prototype components from interface-level or future components. Only measured components contribute to the numerical performance claims in Section 5.

Table 3. Implementation and evaluation status of framework components

Component	Status in this study	Evidence / metric	Production requirement
Fabric ledger, hashes, contracts, audit index	Implemented prototype	Latency, TPS, functional tests, retrieval time	Multi-orderer deployment and Caliper benchmarking
Off-chain ciphertext and access-policy service	Implemented prototype	Hash, ABE, PRE, authorization and interception tests	HSM-backed keys and interdepartmental PKI
Eligibility proof service	Implemented prototype at proof-of-knowledge / predicate level	End-to-end timing and 2048-bit arithmetic microbenchmark	Audited curve/SNARK library and formal circuit review
Federated learning and secure aggregation	Implemented horizontal-FL simulation	Logistic-regression accuracy, non-IID, poisoning, secure-aggregation overhead	Authorized regional application data and secure execution

Group-signature interface	Interface-level design	No latency claim	Audited dynamic membership and opening protocol
Homomorphic encryption	Conceptual optional module	No end-to-end benchmark	Scheme selection according to supported operations and security level

3. System Model and Overall Framework

3.1 Overall Design Objectives

To address data silos, privacy leakage, manual-review burden, fraudulent applications, and untraceable business processes, this paper constructs a trusted data-sharing framework integrating a consortium blockchain with implemented privacy services. The measured prototype includes on-chain/off-chain evidence storage, zero-knowledge eligibility proofs, attribute-based encryption, proxy re-encryption, secure aggregation, horizontal federated logistic regression, and smart-contract auditing. Homomorphic encryption and a future group-signature service with dynamic membership/revocation support are retained only as optional interface-level extensions and are excluded from the numerical claims.

The proposed framework follows the design principles of "on-chain trustworthiness, off-chain security, computational privacy, controllable access, and process traceability". On-chain components mainly record data digests, authorization

credentials, review status, fund flows, and audit logs; off-chain components store raw business data and encrypted materials; the privacy computing module is responsible for eligibility verification, ciphertext sharing, collaborative modeling, and identity protection; and smart contracts are responsible for automatically executing business processes such as application acceptance, permission authorization, eligibility verification, risk warning, fund disbursement, and audit traceability. Through this design, the system can achieve secure collaboration and trusted supervision for childcare subsidy services without breaking departmental data boundaries.

3.2 Overall System Architecture

The childcare subsidy blockchain-based privacy protection and trusted sharing framework designed in this paper consists of four layers: business participation, data storage, privacy computing, and blockchain trusted collaboration. The overall architecture is shown in Fig. 3.




Figure 3. Overall architecture of blockchain-based privacy protection and trusted data sharing for the smart civil-affairs childcare subsidy system

The business participation layer includes eight major entity types: applicant families, the lead subsidy authority, public security, health, human-resources/social-insurance, finance, payment institutions, and audit bodies. The data layer combines departmental source systems, off-chain ciphertext storage, and on-chain hashes and indexes. The measured privacy layer includes zero-knowledge proof, attribute-based encryption, proxy re-encryption, Fabric MSP/CA certificate authentication with pseudonymous identifiers, and horizontal federated learning across regional subsidy-service nodes. Homomorphic encryption and a future group-signature service with dynamic membership/revocation support are optional interfaces and are not benchmarked. The blockchain layer records evidence, authorization decisions, workflow states, model-update digests, disbursement events, and audit records.

3.3 Formal System Model

To facilitate the description of multi-entity collaboration relationships in the proposed framework, the system participant set is defined as:

$$P = \{P_f, P_m, P_p, P_h, P_s, P_i, P_b, P_a\} \quad (7)$$

where Pf denotes the applicant family, Pm denotes the lead subsidy authority, Pp denotes the public security department, Ph denotes the health commission, Ps denotes the human resources and social security/medical insurance department, Pi denotes the finance department, Pb denotes the financial institution, and Pa denotes the regulatory audit department. Different entities participate in application, verification, review, disbursement, and auditing according to their business responsibilities, and form trusted collaboration relationships in the consortium blockchain.

3.4 Threat Model and Security Assumptions

The revised threat model includes semi-honest participants, malicious external attackers, malicious or colluding departmental nodes, compromised organizational certificates or data keys, replayed authorization and verification requests, buggy or malicious smart contracts, federated-learning poisoning, and model-update inference. It also considers a governance conflict between immutable audit evidence and the correction, retention, revocation, or deletion of minors' off-chain records. Attackers may know the system architecture and public policies but are assumed not to break standard cryptographic primitives or compromise a threshold of consensus and audit authorities simultaneously. The corresponding threats and protection mechanisms are summarized in Table 4.

Table 4. System Threats and Protection Mechanisms

Threat Type	Risk Description	Protection Mechanism
Data Leakage	Sensitive data such as birth medical certificates, household registration information, and bank account details are illegally obtained	Off-chain ciphertext storage, attribute-based encryption, proxy re-encryption
Identity Association	Attackers link the identity of applicant households through on-chain transactions or access records	MSP/CA certificates; pseudonymous IDs; minimized on-chain metadata; optional group-signature interface
Unauthorized Access	Internal personnel access sensitive data beyond the scope of their job responsibilities	Dynamic access control, authorization contract, audit contract
Data Tampering	Application materials, review status, or fund disbursement records are maliciously modified	Blockchain deposit, digital signature, timestamp
Subsidy Fraud	Behaviors including duplicate application, cross-region illegal collection, submission of false materials, or identity impersonation	Zero-knowledge proof, federated learning, secure aggregation
Model Inference Attack	Attackers infer the data characteristics of other departments through the uploaded model parameters	Secure aggregation, parameter summary on-chain, abnormal node audit
Node Collusion / FL Poisoning	One or more regional FL clients submit coordinated sign-flip or model-replacement updates	Update clipping, validation-score monitoring, trimmed mean/coordinate median, on-chain accountability

Compromised Keys or Certificates	A stolen departmental credential is used for unauthorized access or forged requests	Short-lived certificates, key rotation, revocation list, HSM, emergency suspension
Replay Attack	A valid request, proof, or authorization token is resubmitted	Nonce, timestamp, application-state check, token expiry, replay cache
Malicious or Vulnerable Smart Contract	A contract contains permission bypass, unsafe upgrade, or inconsistent state transition	Multi-party code review, versioned endorsement, least privilege, rollback and emergency pause
Data-Lifecycle Conflict	A minor-related record must be corrected, expired, revoked, or deleted although audit evidence is immutable	Off-chain correction/deletion, superseding on-chain event, key destruction, retention policy and legal audit trail

Security claims are conditional on an honest majority of consortium governance authorities, protected local databases, secure key generation and storage, audited contract code, and legally authorized tracing. The framework does not protect against simultaneous compromise of the lead authority, a

consensus majority, and the audit authority. The added poisoning tests cover up to two colluding clients among five participants; stronger Byzantine settings require threshold governance and more advanced robust aggregation.

3.5 Framework Operation Process




Figure 4. Operational workflow of the smart civil-affairs childcare subsidy system

The operation process of the proposed framework is shown in Fig. 4 and contains seven system-processing steps: application submission, material evidence storage,

authorization verification, eligibility review, risk identification, fund disbursement, and audit traceability. The applicant submits the application and materials through the

service platform. The measured prototype authenticates organizations through Fabric MSP/CA certificates and represents applicant cases through pseudonymous application identifiers; a group-signature interface is optional and is not used in the reported tests. Materials are encrypted off chain, while the ciphertext hash, storage index, timestamp, and submitter signature are written to Fabric. The lead subsidy authority initiates authorized verification requests to the supporting departments. Verification departments return zero-knowledge proof results or verification digests rather than raw records. A horizontal federated risk model then operates across regional subsidy-service nodes that hold disjoint applicant records under a common feature schema. High-risk cases enter manual review; approved cases trigger disbursement and receipt confirmation. Finally, the audit contract records the complete event chain.

3.6 Data Correction, Retention, Revocation, and Deployment Governance

Blockchain immutability is applied only to evidence and event history, not to perpetual storage of minors' raw records. When source data are corrected, the off-chain record is replaced under an authorized workflow and a superseding on-chain event links the old and new hashes without overwriting history. When retention expires or deletion is legally required, the ciphertext and associated data-encryption key are deleted or cryptographically destroyed; the ledger retains only a minimal, non-content-bearing audit event. Certificate revocation, emergency access, policy-version rollback, and key rotation are governed by multi-party contracts. Deployment should establish an interdepartmental data dictionary, certificate and key lifecycle governance, contract-upgrade procedures, incident-response procedures, and clear organizational authorization boundaries; related public-sector governance and broader blockchain-adoption considerations are discussed in [10-12]. Any processing or sharing of protected personal data should additionally comply with applicable legal and regulatory requirements.

4. Privacy Protection and Trusted Data Sharing Mechanism

4.1 Overall Mechanism Design

Based on the system model and overall framework in Chapter 3, this chapter further designs a privacy protection and trusted data sharing mechanism for smart civil affairs childcare subsidy systems. The mechanism focuses on key stages such as application material submission, cross-departmental eligibility verification, dynamic access control, fraudulent application identification, fund disbursement, and audit traceability. Its goal is to achieve trusted sharing of subsidy data, privacy-preserving eligibility verification, dynamic access permission management, intelligent identification of

fraudulent applications, and traceable business process supervision without centrally aggregating raw sensitive data.

The proposed mechanism consists of five parts: on-chain/off-chain trusted storage, dynamic access control and ciphertext authorization, zero-knowledge eligibility privacy verification, federated learning-based fraudulent application identification, and smart contract auditing. Among them, the on-chain/off-chain trusted storage mechanism verifies data integrity without uploading raw sensitive data to the blockchain; the dynamic access control and ciphertext authorization mechanism ensures controllable data access permissions across different departments, positions, and business stages; the zero-knowledge eligibility privacy verification mechanism completes subsidy eligibility verification without disclosing complete birth, household registration, or guardianship relationship data; the federated learning-based fraudulent application identification mechanism identifies duplicate applications, cross-regional illegal claims, and abnormal fund flows without moving raw data out of local domains; and the smart contract audit mechanism records the full process of application, verification, review, disbursement, and supervision to form a tamper-proof evidence chain.

4.2 On-Chain/Off-Chain Trusted Storage and Data Integrity Verification

Childcare subsidy services involve multiple types of sensitive data, including identity certificates, birth medical certificates, household registration relationships, guardianship relationships, bank accounts, historical subsidy records, and review opinions. Directly storing raw data on the blockchain would not only impose a heavy storage burden, but also cause privacy leakage risks due to the transparency and immutability of the blockchain ledger. Therefore, this paper adopts on-chain/off-chain collaborative storage. Raw data, certificates, and business materials are encrypted and stored in departmental local databases, government cloud systems, or off-chain systems such as IPFS, while only data hashes, ciphertext indexes, metadata, timestamps, and signatures are stored on-chain.

Suppose that the i -th business data item submitted by an applicant family is d_i . The system generates a data key k_i according to its sensitivity level and encrypts the raw data:

$$C_i = \text{Enc}_{k_i}(d_i) \quad (8)$$

where C_i denotes the off-chain ciphertext data, and $\text{Enc}_{k_i}(\cdot)$ denotes the encryption algorithm based on key k_i . The system further calculates the ciphertext hash:

$$h_i = \text{Hash}(C_i) \quad (9)$$

The corresponding on-chain evidence record is defined as:

$$R_i = \{h_i, \text{CID}_i, \text{Meta}_i, \text{Time}_i, \text{Sig}_i\} \quad (10)$$

where CIDI denotes the off-chain storage index, Metai denotes metadata such as data type, application number, sensitivity level, and business stage, Timei denotes the timestamp, and Sigi denotes the signature of the submitting entity. When any authorized entity later accesses off-chain data, the system recalculates the ciphertext hash and compares it with the on-chain record hi. If the two are consistent, the off-chain data have not been tampered with; otherwise, an integrity risk is identified. This mechanism ensures the integrity and traceability of application materials, review data, and disbursement records while avoiding direct uploading of raw data to the blockchain.

4.3 Dynamic Access Control and Ciphertext Authorization Sharing

Data access permissions in the childcare subsidy system are not fixed, but are closely related to departmental responsibilities, job roles, business stages, data categories, administrative regions, and authorization periods. To avoid unauthorized access and delayed permission revocation caused by traditional static role authorization, this paper constructs a dynamic access control mechanism based on attribute-based encryption and combines it with proxy re-encryption for cross-departmental ciphertext authorization sharing, drawing on established ABE, PRE, and multi-layer access-control research [32-34,44]. Let the attribute set of access subject u be:

$$\text{Attr}(u) = \{a_1, a_2, \dots, a_n\} \quad (11)$$

where the attributes may include institution type, job responsibility, business permission, administrative region, authorization period, and security level. Let Policy(d,t) denote the access policy of data d at business stage t; then the access decision function is defined as Eq. (5).

To support cross-departmental ciphertext forwarding, this paper introduces a proxy re-encryption mechanism [34]. When data owner Px authorizes data user Py to access ciphertext data, the system generates a re-encryption key $\text{rk}_{x \rightarrow y}$, and the proxy converts ciphertext Cx into Cy without obtaining the plaintext:

$$\text{ReEnc}(C_x, \text{rk}_{x \rightarrow y}) \rightarrow C_y \quad (12)$$

This process avoids plaintext transmission and repeated decryption of raw data, improving the security and flexibility of cross-departmental data sharing. When a position changes, authorization expires, or risk status changes, the authorization contract can update the access policy and prevent invalid subjects from continuing to access historical data through key update, revocation lists, or re-encryption strategies.

4.4 Zero-Knowledge Proof-Based Eligibility Privacy Verification

Childcare subsidy eligibility verification usually requires verifying the applicant child's age, birth information, household registration status, guardianship relationship, and whether duplicate claims exist. Traditional verification methods require transmitting a large amount of raw data between departments, which can easily lead to privacy leakage. Following established zero-knowledge proof principles [36,37] and their use in privacy-preserving blockchain systems [35], this paper designs an eligibility privacy verification mechanism in which eligibility can be verified without exposing the underlying sensitive attributes.

Let S denote the subsidy eligibility proposition, x denote public input, and w denote private witness. Here, x may include the application number, policy conditions, data hash, and verification time, while w may include sensitive information such as birth date, guardianship relationship, household registration status, and historical claim records. The zero-knowledge proof process can be expressed as:

$$\text{Prove}(x, w) \rightarrow \pi, \quad \text{Verify}(x, \pi) \rightarrow \{0, 1\} \quad (13)$$

where π denotes the proof result. If the verification result is valid, the applicant satisfies the corresponding eligibility conditions; otherwise, the proof is invalid or the eligibility conditions are not met. In this process, the verifier obtains only whether the proof holds and cannot access the specific content of private witness w.

To meet the multi-condition verification requirements of childcare subsidies, the eligibility proposition is divided into m sub-propositions:

$$S = S_1 \wedge S_2 \wedge \dots \wedge S_m \quad (14)$$

where S1 indicates that the child's age satisfies the policy requirement, S2 indicates that the applicant has a legal guardianship relationship with the child, S3 indicates that household registration or residence status satisfies regional policy conditions, and S4 indicates that no duplicate claim record exists. The overall eligibility decision function is:

$$\text{Eligible} = \begin{cases} 1, & \bigwedge_{r=1}^m \text{Verify}(x_r, \pi_r) = 1 \\ 0, & \text{otherwise} \end{cases} \quad (15)$$

where π_r denotes the proof corresponding to the r-th sub-proposition. If all sub-proofs pass verification, the system determines that the applicant is eligible for the subsidy; otherwise, the process enters material supplementation, manual review, or rejection. This mechanism significantly reduces the exposure scope of raw data in cross-departmental verification, so that the health commission, public security, and human resources and social security/medical insurance departments submit only proof results or verification digests, rather than complete raw records to the lead subsidy authority.

4.4.1 Lightweight Proof Optimization

The proof service separates offline and online work. Random commitments for the four eligibility sub-propositions are

generated during idle time and stored in a one-time commitment pool. At request time, the service binds the precomputed commitments to the application hash, policy version, and nonce, computes the challenge and response, and immediately erases the consumed randomness. Proofs are generated off chain and only the verification result and proof digest are written to Fabric. The optimization reduces online modular exponentiation while preserving freshness through nonces and one-time commitment use. Section 5.3.2 reports both end-to-end service overhead and a reproducible 2048-bit arithmetic microbenchmark.

4.5 Federated Learning-Based Fraudulent Application Identification

Fraud risks in childcare-subsidy services have cross-departmental and cross-regional correlations. Departmental systems first verify their responsibility-specific fields locally and return authorized proof results or status indicators. Each regional subsidy-service node then constructs the same risk-feature vector for its own disjoint set of applications. The federated-learning experiment is therefore horizontal rather than vertical: clients share a common model and feature schema but do not share applicant records. This design supports regional joint risk identification while preserving departmental source-data boundaries.

Suppose that there are K participating regional subsidy-service nodes, each holding a disjoint local applicant dataset with the same risk-feature schema. The k -th node contains n_k samples, and the total number of samples is n . The global optimization objective follows a standard horizontal federated-learning formulation [18-21,30]:

$$\min_w F(w) = \sum_{k=1}^K \frac{n_k}{n} F_k(w) \quad (16)$$

where w denotes the shared logistic-regression parameter and $F_k(w)$ denotes the local loss of the k -th regional node. Each client trains the same model locally and uploads only its update. The coordinator uses secure aggregation with input validation and dropout-aware controls [27,28], rather than receiving applicant-level records:

$$\Delta w = \sum_{k=1}^K \frac{n_k}{n} \Delta w_k \quad (17)$$

where Δw_k denotes the local model update of the k -th node and Δw denotes the aggregated global model update. To enhance the trustworthiness of the federated learning process, this paper writes training summaries to the blockchain. The training record of the r -th round is defined as:

$$FL_r = \{\text{Round}_r, \text{NodeSet}_r, \text{Hash}(\Delta w_r), \text{Score}_r, \text{Time}_r, \text{Sig}_r\} \quad (18)$$

where Round_r denotes the training round, NodeSet_r denotes the set of participating nodes, $\text{Hash}(\Delta w_r)$ denotes the model update digest, Scorer denotes the model validation score or risk identification performance indicator, Timer denotes the record time, and Sigr denotes the submitter's signature. Through this structure, the system can record the entire federated training process, prevent model updates from being tampered with, and support abnormal node accountability.

For the i -th applicant, the system generates a risk score based on the trained model:

$$\text{Risk}_i = f(x_i, w) \quad (19)$$

where x_i denotes the risk features of the applicant, w denotes the trained model parameters, and Risk_i denotes the risk score. When $\text{Risk}_i \geq 0$, the system determines that the application has a high risk and enters manual review or audit; when $\text{Risk}_i < 0$, the application proceeds to the normal review process. This mechanism enables cross-departmental joint risk control without moving raw data out of local domains, thereby improving the ability to identify subsidy fraud.

4.5.1 Non-IID and Collusion-Resistant Aggregation

Before aggregation, each regional update is checked for norm, direction consistency, and validation loss. Updates exceeding the clipping radius are scaled to the radius. If all updates pass, weighted FedAvg is used. When one suspicious update is detected among five clients, a coordinate-wise trimmed mean removes the largest and smallest coordinate values; when two coordinated suspicious updates are detected, the coordinate median is used. The chain records the update digest, validation score, clipping decision, and aggregation rule. The input-validation and dropout-resistance objectives are consistent with recent privacy-preserving secure-aggregation research [28]. This mechanism is a prototype defense rather than a proof of Byzantine robustness.

4.6 Smart Contract Auditing and Core Process

To achieve trusted full-process supervision of childcare-subsidy services, the framework uses application, authorization, verification, risk-control, disbursement, and audit contracts. Smart-contract-based access control can enforce static and dynamic authorization rules through distributed contracts [9]. The application contract records the application number, material hash, and state; the authorization contract evaluates attributes and versioned policies; the verification contract checks proof results or digests; the risk-control contract consumes the federated risk score; the disbursement contract records allocation and receipt status; and the audit contract links the complete event sequence. For any business event e_j , its on-chain audit record is defined as:

$$\text{Audit}_j = \{ \text{EventID}_j, \text{ApplID}_j, \text{Actor}_j, \text{Action}_j, \text{Result}_j, \text{Time}_j, \text{Sig}_j \} \quad (20)$$

where EventID_j denotes the event number, ApplID_j denotes the application number, Actor_j denotes the operating subject, Action_j denotes the operation type, Result_j denotes the operation result, Time_j denotes the event occurrence time, and Sig_j denotes the signature of the operating subject. Through this structure, the system can trace the complete process of any application from submission, verification, and review to disbursement and audit.

To further describe the overall execution process of the proposed mechanism, this paper summarizes the privacy-preserving eligibility verification and trusted sharing process as Algorithm 1.

Algorithm 1 summarizes the privacy-preserving eligibility-verification and trusted-sharing procedure. It integrates off-chain encryption, Fabric evidence records, attribute-based authorization, zero-knowledge eligibility proofs, horizontal federated risk identification, disbursement control, audit logging, and lifecycle rules into one standard pseudocode block.

Algorithm 1. Privacy-Preserving Eligibility Verification and Trusted Sharing (standard pseudocode)

Item	Procedure
Input	Application data d _i , requester u, Attr(u), Policy(d,t), eligibility predicates S ₁ ...S _m , risk threshold theta.
Output	Eligibility result, risk score, disbursement status, and audit record.
1	Authenticate the requester and create a unique application and nonce.
2	Encrypt d _i off chain; compute h _i = Hash(C _i); write the minimal evidence record to Fabric.
3	Evaluate Access(u,d,t); reject, log, and stop if the policy is not satisfied.
4	Authorize ciphertext access using ABE/PRE and a versioned, expiring authorization token.
5	Generate one-time predicate proofs from local witnesses and verify all subproofs on the verification service.
6	If eligibility fails, write the reason code and route to supplementation/manual review.
7	Train departmental risk models locally; clip and securely aggregate validated updates.
8	Compute the risk score; route high-risk cases to manual audit.
9	For approved cases, invoke the disbursement contract and record payment/receipt status.
10	Write the complete event chain; apply correction, revocation, and retention rules when required.

4.7 Practical Deployment Roadmap

A staged deployment is recommended. Stage 1 establishes the interdepartmental data dictionary, legal authorization matrix, PKI/HSM service, and one-way audit integration with existing systems. Stage 2 pilots on-chain evidence, authorization, and audit contracts in one district while retaining existing approval and payment systems as systems of record. Stage 3 introduces minimum-disclosure proofs and federated risk detection across several districts. Stage 4 establishes a cross-region consortium with at least three Orderer nodes, policy-version governance, interoperability gateways, disaster recovery, and independent security audits. This roadmap separates technical feasibility from institutional authorization and reduces migration risk.

5. Experimental Results and Performance Analysis

5.1 Experimental Environment and Experimental Scheme

5.1.1 Experimental Environment and Platform Configuration

The prototype uses Hyperledger Fabric 2.4 with seven blockchain organizations representing the lead subsidy authority, public security, health, human-resources/social-insurance, finance, a payment institution, and audit. Each organization hosts one Peer and organization-specific MSP credentials. A single application channel is used, TLS is enabled, and the endorsement policy requires any three

authorized organizations relevant to the invoked contract. The development environment uses one Raft Orderer and a 1 s batch interval; production deployment requires at least three Orderers. The blockchain organizations are not identical to the five federated-learning clients: the reported horizontal FL experiments use five regional subsidy-service nodes with a

common feature schema and disjoint applicant records. Fig. 10 contains point estimates because repeated-run Fabric logs were not preserved; the added proof, non-IID, and poisoning tests report repeated-run statistics. The detailed configuration is summarized in Table 5.

Table 5. Experimental Environment Configuration

Category	Configuration Details
CPU	Intel Xeon Silver 4210R, 2.40 GHz
Memory	32 GB
Operating System	Ubuntu 20.04 LTS
Blockchain Platform	Hyperledger Fabric 2.4
Container Environment	Docker 20.10, Docker Compose 1.29
Smart Contract Language	Go 1.19
Off-chain Storage	IPFS 0.18, Local MySQL 8.0
Privacy Computing Environment	Python 3.9, PyTorch 1.13
Federated Learning Model	Logistic regression (all reported Figs. 13-16)
Consensus Mechanism	Raft Ordering Service
Block Size	100 transactions per block
Transaction Timeout	2 s
Network Topology	7 organizations, 7 Peer nodes, 1 development Raft Orderer, 1 application channel
Endorsement Policy	Any 3 authorized organizations required; contract-specific organization checks
Identity and Transport	Fabric MSP / CA credentials, TLS enabled, certificate revocation list
Block Generation	100 transactions per block or 1 s batch interval, whichever occurs first
Reporting Protocol	Original Fabric performance: point estimates; added FL: 5 runs; proof arithmetic: 100 requests

The Fabric network simulates multi-organization government collaboration. Each blockchain organization deploys one Peer; the Orderer performs transaction ordering and block generation. Chaincode implements evidence, authorization, verification, disbursement, and audit events. Off-chain IPFS stores encrypted materials, and local databases simulate departmental sources. For fraud-risk modeling, verified status indicators are transformed locally into a common risk-feature schema and assigned to regional subsidy-service clients by application region; this separation resolves the distinction between departmental verification roles and horizontal FL clients.

5.1.2 Experimental Dataset and Business Scenario Construction

No identifiable Chinese childcare-subsidy microdata were used because the authors did not possess data-use authorization covering minors' records. The experimental dataset is explicitly synthetic. Records are generated from the fields in Table 6 using bounded categorical and Bernoulli distributions. Normal applications account for 85%; the remaining 15% are divided equally among duplicate application, cross-regional duplicate claim, identity abnormality, guardianship abnormality, and bank-account abnormality. Correlated rules ensure that abnormal labels are linked to the relevant risk fields.

Table 6. Simulated Childcare Subsidy Application Data Fields

Field Category	Field Name	Field Description
Application Information	ApplID	Application serial number
Identity Information	ApplicantIDHash	Hash value of the applicant's identity
Infant Information	ChildIDHash, ChildAge	Hash value of the infant's identity and age
Birth Information	BirthCertStatus	Status of the birth medical certificate
Household Registration Information	HouseholdStatus	Household registration status
Guardianship Relationship	GuardianRelation	Validity of the guardianship relationship
Social Security & Medical Insurance	SocialInsuranceStatus	Social security/medical insurance status
Historical Subsidy Records	SubsidyHistory	Existence of historical subsidy collection records
Bank Account	BankAccountStatus	Validity of the bank account
Regional Information	RegionCode	Code of the application region
Risk Label	RiskLabel	Normal/abnormal application label

Four scales (5,000, 10,000, 20,000, and 50,000 records) are generated. Records are divided into training, validation, and test sets at 70:15:15 by application identifier to prevent duplicate-record leakage. Random seed 42 is used for the baseline, and seeds 42-46 are used for the five-run robustness tests. For horizontal FL, every regional client receives the same non-identifier risk-feature schema and a disjoint set of applications. Near-IID clients are randomly partitioned by application; moderate and strong non-IID clients are created by skewing regional risk prevalence and selected feature distributions. Departmental systems remain source-data verifiers and are not modeled as clients with incompatible feature spaces.

5.1.3 Literature-Grounded Workflow and Peer-Reviewed Administrative-Data Validation

To strengthen external validity without overstating access to protected records, the revision performs two literature-based checks. First, the workflow is compared with peer-reviewed evidence on eligibility determination and redetermination, reporting burden, participation continuity, and subsidy administration [1,2]. Second, the data model is compared with the administrative-data evaluation framework proposed by Horak and Ward [3]. This is a literature-grounded schema-and-process validation rather than a fraud-detection accuracy test on real government microdata. The evidence in [1-3] is used only for process-level and data-model comparison; it is not treated as evidence that agency assignments or legal rules are identical across jurisdictions. Because the current synthetic schema does not include complete provider- and service-level records, the result is reported as partial structural alignment. The results are summarized in Table 7.

Table 7. Literature-grounded workflow and administrative-data validation

Peer-reviewed evidence / data element	Framework representation	Validation outcome
Eligibility determination and periodic redetermination [1]	Application contract; versioned eligibility policy; authorization and state records	Core eligibility workflow represented
Reporting requirements and administrative burden [1]	Minimal-field requests; reusable verification results; exception and renewal states	Burden-reduction mechanisms represented at workflow level

Participation continuity, duration, and geographic variation [1,2]	SubsidyHistory; RegionCode; status transitions; longitudinal audit events	Longitudinal participation monitoring supported
Child, family, eligibility, and subsidy-participation information [2,3]	ChildIDHash; ApplicantIDHash; eligibility predicates; bank/payment and benefit-history entities	Core administrative categories represented
Provider-level administrative data [3]	No complete provider/service table in the current synthetic schema	Partial structural alignment; schema extension required
Program evaluation, data quality, and performance monitoring [3]	Risk-control contract; traceability index; manual-review route; audit records	Supports data-driven evaluation and accountable exception handling

5.1.4 Comparison Schemes and Evaluation Metrics

The evaluation uses three internal baselines: a centralized database with rule/manual review, a basic blockchain evidence scheme, and a single-encryption sharing scheme. These controlled baselines isolate the effects of ledger

evidence, privacy mechanisms, collaborative risk detection, and auditing. Their functional characteristics are illustrated in Fig. 5. Table 1 separately compares functional coverage with published studies [7-9,13,14,23,24,44]; published numerical results are not mixed because their datasets and platforms differ.




Figure 5. Functional comparison of the internal data-sharing schemes

This paper selects security, transaction latency, throughput, encryption and decryption time, zero-knowledge proof time, federated learning model performance, and traceability efficiency as evaluation metrics. Block transaction confirmation latency is defined as:

$$T_{tx} = T_{confirm} - T_{submit} \quad (21)$$

where T_{submit} denotes the transaction submission time and $T_{confirm}$ denotes the transaction confirmation time. System throughput is defined as:

$$TPS = \frac{N_{tx}}{T_{total}} \quad (22)$$

where N_{tx} denotes the number of successfully confirmed transactions and T_{total} denotes the total processing time. Federated learning risk identification performance is measured by accuracy, precision, recall, and F1 score:

$$Acc = \frac{TP + TN}{TP + TN + FP + FN} \quad (23)$$

$$Precision = \frac{TP}{TP + FP} \quad (24)$$

$$Recall = \frac{TP}{TP + FN} \quad (25)$$

$$F1 = \frac{2 \times Precision \times Recall}{Precision + Recall} \quad (26)$$

where TP, TN, FP, and FN denote true positives, true negatives, false positives, and false negatives, respectively.

5.2 Functional Testing

5.2.1 Identity Authentication and Dynamic Permission Control Test

The identity authentication and permission control test is used to verify whether the system can dynamically determine access permissions according to subject identity, departmental responsibility, position attributes, and business stage. The experiment sets up six types of access subjects, including civil affairs, public security, health commission, finance, financial institution, and audit department, and constructs legitimate access requests, unauthorized access requests, requests after authorization expiration, and requests after permission revocation. The system reads the subject attribute set and access policy through the authorization contract, determines whether the subject has access permission, and writes the access result into the audit log. The test results are shown in Fig. 6.




Figure 6. Functional results of identity authentication and dynamic access control

As shown in Fig. 6, the system can correctly identify legitimate and unauthorized requests according to the access policy and block access after permission expiration or revocation, indicating that the dynamic access control

mechanism can meet the multi-departmental, multi-stage, and classified authorization requirements of childcare subsidy services.

5.2.2 Cross-Departmental Privacy-Preserving Data Sharing Test

The cross-departmental data sharing test is used to verify the system's ability to complete eligibility verification without exposing raw sensitive data. The experiment simulates verification requests initiated by the lead subsidy authority to the public security, health commission, and human resources and social security/medical insurance departments. The public security department is responsible for identity, household registration, and guardianship relationship verification; the health commission department is responsible for birth medical certificate verification; and the human resources and social security/medical insurance department is

responsible for benefit status and historical claim record verification. The system uses off-chain ciphertext storage, attribute-based encryption, proxy re-encryption, and zero-knowledge proof to complete sharing and verification.

The experiment sets 1,000 cross-departmental verification requests, including 800 legitimate requests and 200 illegal or unauthorized requests. The 300 birth-certificate tests, 300 household/guardianship tests, and 200 social-insurance/benefit-status tests form mutually exclusive subsets of the 800 legitimate requests. Test indicators include authorization success, illegal-request interception, raw-record exposure, verification accuracy, and audit-record completeness. The results are shown in Fig. 7.




Figure 7. Cross-departmental privacy-preserving data-sharing test results

As shown in Fig. 7, the system can effectively support privacy-preserving data sharing among multiple departments. The authorization success rate of legitimate cross-departmental verification requests reaches 99.50%, and the interception rate of illegal or unauthorized requests reaches 100%, indicating that the dynamic access control mechanism can accurately distinguish legal access from unauthorized access. During verification of sensitive data such as birth medical certificates, household registration and guardianship relationships, and social security/medical insurance status, the system returns only verification conclusions, proof digests, or zero-knowledge proof verification results, without exposing raw sensitive data in plaintext. Meanwhile, all cross-departmental requests generate on-chain audit records, demonstrating that the proposed scheme ensures traceability of data sharing while protecting privacy.

5.2.3 Smart Contract Automatic Review and Traceability Audit Test

The smart-contract review and traceability test simulates 1,000 applications, including 850 baseline-normal and 150 risk-labeled applications. Within the 150 risk-labeled cases, 80 are duplicate/cross-regional claims and 70 are high-risk cases routed to manual review; these two groups are non-overlapping. Material-integrity, eligibility-condition, account-status, disbursement, and audit-completeness checks are rule-level functional tests evaluated across the full 1,000 applications and may overlap the normal/risk labels. This distinction prevents the plotted test items from being interpreted as mutually exclusive class counts. The results are shown in Fig. 8.




Figure 8. Smart-contract automatic review and traceability test results

As shown in Fig. 8, the smart contracts designed in this paper can effectively support automatic review and full-process auditing of childcare subsidy services. The system achieves high accuracy in material integrity verification, eligibility condition judgment, duplicate application interception, and transfer of high-risk applications to manual review. In particular, the audit log generation integrity rate and abnormal application traceability success rate both reach 100%, indicating that key events such as application, authorization, verification, risk control, disbursement, and audit can be completely recorded. This result shows that the

smart contract collaboration mechanism can effectively improve the automation level, review consistency, and regulatory traceability of childcare subsidy business processes.

5.2.4 Quantitative Security Evidence by Component

Figure 9 consolidates the security evidence that is directly measured in the prototype. It distinguishes observed indicators from conceptual security properties and avoids treating the presence of a cryptographic term as experimental proof of security.




Figure 9. Quantitative security evidence for implemented components

5.3 Performance Test and Comparative Analysis

5.3.1 Block Transaction Latency and Throughput Test

Block transaction latency and throughput are feasibility indicators for the stated Fabric configuration. Concurrent

requests are set to 50, 100, 200, 400, and 800. Because raw repeated-run logs from the original blockchain benchmark were not retained, the values in Fig. 10 are explicitly reported as point estimates rather than mean $\pm$ standard deviation. The revision therefore narrows the claim to feasibility under the controlled configuration and recommends repeated Hyperledger Caliper tests in future deployment work.




Figure 10. Block transaction latency and throughput under different concurrency levels

As shown in Fig. 10, when the number of concurrent requests increases from 50 to 400, the system throughput increases from 392 TPS to 718 TPS, indicating that system resource utilization gradually improves. When the number of concurrent requests reaches 800, system throughput decreases to 684 TPS and the average transaction latency increases to 329 ms, indicating that the system enters a high-load state. Overall, the proposed scheme maintains high throughput and acceptable transaction confirmation latency under medium concurrency, making it suitable for business scenarios such as batch childcare subsidy applications, centralized review, and periodic disbursement.

5.3.2 Encryption, Decryption, and Zero-Knowledge Proof Time Test

To evaluate the computational overhead introduced by privacy protection mechanisms, the experiment tests the average time consumption of attribute-based encryption, attribute-based decryption, proxy re-encryption, zero-knowledge proof generation, zero-knowledge proof verification, and hash evidence storage. The measured operation-level overhead is summarized in Fig. 11.




Figure 11. Average time consumption of privacy-protection operations

As shown in Fig. 11, hash evidence storage has the lowest overhead and can meet the requirements of high-frequency application material evidence storage. Attribute-based encryption and decryption take 58 ms and 46 ms, respectively, making them suitable for multi-departmental data access control. Proxy re-encryption takes an average of 67 ms and can support cross-departmental ciphertext forwarding. Zero-knowledge proof generation has a relatively high time cost of 143 ms, but proof generation is mainly completed locally by the data-providing department and does not significantly affect on-chain review efficiency. Zero-knowledge proof verification takes 86 ms, satisfying the combined security and response efficiency requirements of childcare subsidy eligibility verification.

5.3.2.1 Lightweight Proof Precomputation Microbenchmark

A separate Python 3.9 arithmetic microbenchmark evaluates four Schnorr-style proof-of-knowledge subproofs per request using a 2048-bit MODP group and SHA-256, following the classical Schnorr identification/proof framework [45]. One hundred requests are repeated. This microbenchmark isolates cryptographic arithmetic and is not directly comparable with the 143 ms end-to-end proof-service time in Fig. 11, which includes serialization, policy evaluation, inter-process communication, and persistence. The results are reported in Fig. 12.




Figure 12. Lightweight proof arithmetic with offline commitment precomputation

The optimization moves commitment generation outside the interactive request path. Commitments are consumed once and bound to a nonce and policy version. The result supports a lightweight online path, but production use requires an audited proof library, secure randomness, and formal verification of the eligibility circuit.

5.3.3 Federated Learning Training and Aggregation Efficiency Test

To evaluate horizontal federated learning for fraudulent-application identification, the experiment uses five regional

subsidy-service clients. Every client holds disjoint application records represented by the same feature schema. All reported results in Figs. 13-16 use logistic regression with the same parameter vector; a lightweight MLP is retained only as an unreported implementation option. Data scales are 5,000, 10,000, 20,000, and 50,000 records. Secure aggregation is performed after each round, and the model digest is written to the blockchain. Training efficiency is summarized in Fig. 13.




Figure 13. Federated-learning training-time breakdown under different data scales

As shown in Fig. 13, the total time of federated learning increases with data scale, and local training time is the main source of overhead. Secure aggregation and summary on-chain time are relatively low; under the scale of 50,000 records, they are 0.83 s and 0.34 s, respectively. This indicates that recording federated learning summaries on the blockchain does not impose a significant burden on overall training efficiency. The results show that the proposed

scheme can achieve acceptable cross-departmental joint risk control training on medium-scale childcare subsidy datasets.

The deployment-mode comparison uses the same logistic-regression risk classifier for the single-node, centralized, federated, and federated-plus-secure-aggregation settings. The results are shown in Fig. 14.




Figure 14. Performance comparison of logistic-regression risk models under different deployment modes

As shown in Fig. 14, the single-department local model has limited data dimensions and is difficult to use for identifying cross-departmental correlated risks, resulting in an accuracy of only 84.5%. The centralized model can directly use the full dataset and reaches an accuracy of 94.1%, but it requires the aggregation of raw sensitive data and is therefore unsuitable for real government privacy scenarios. The federated learning model achieves an accuracy of 91.8% without moving raw data out of local domains, which is close to the centralized model. After secure aggregation is introduced, the model accuracy slightly decreases to 91.1%, but parameter privacy protection is improved, indicating that the proposed scheme achieves a good balance between risk identification effectiveness and privacy protection.

5.3.3.1 Federated Learning under Non-IID Departmental Data

Figure 15 reports a separate robustness experiment on a fixed 30,000-record synthetic dataset and should not be directly compared with the different-scale baseline in Fig. 14. Five horizontal FL clients use a 70:15:15 split and seeds 42-46. Near-IID clients are randomly partitioned by application. Moderate and strong non-IID settings introduce increasingly skewed regional risk prevalence and covariate distributions. Results are reported as mean $\pm$ standard deviation over five runs.




Figure 15. Federated Learning Performance under Different Non-IID Distributions

5.3.3.2 Multi-Node Collusion and Model-Poisoning Test

One or two of five regional clients perform a coordinated sign-flip/model-replacement attack. Standard FedAvg is compared with adaptive norm clipping plus a trimmed mean for one suspicious client or a coordinate median for two suspicious clients. The same fixed dataset, logistic-regression

model, and five random seeds are used. Fig. 16 reports the collusion-resistance results. F1 is calculated within each run and then averaged, so it may differ slightly from an F1 value recomputed from the rounded mean precision and recall shown in the figure.




Figure 16. Resistance to colluding sign-flip clients (five runs)

The attack demonstrates that ordinary FedAvg is vulnerable to coordinated update manipulation. The prototype robust rules recover performance in this controlled five-client setting, but they do not establish general Byzantine security; larger coalitions, adaptive attackers, and backdoor attacks remain future work.

5.3.4 Comprehensive Performance Comparison of Different Schemes

To compare the functional coverage of the proposed scheme with the three internal baselines without introducing unsupported cross-platform numerical claims, the qualitative comparison is shown in Fig. 17. Quantitative latency and throughput are reported only for the stated proposed-system configuration in Fig. 10.

Capability	Scheme A Traditional centralized system	Scheme B Basic blockchain scheme	Scheme C Single-encryption sharing scheme	Proposed Integrated measured framework
Sensitive data kept decentralized	✗ No	✓ Yes	✓ Yes	✓ Yes
Tamper-evident evidence	○ Weak	● Strong	● Medium	● Strong
Fine-grained dynamic access	○ Weak	○ Weak	● Medium	● Strong
Minimum-disclosure eligibility	○ Not supported	○ Not supported	● Partial	● Supported
Joint regional fraud detection	○ Weak	○ Weak	○ Weak	● Strong
Full-process auditability	○ Weak	● Medium	○ Weak	● Strong
Implemented identity protection	Conventional account	Organizational identity	Key authorization	MSP/CA + pseudonymous IDs
Optional interfaces disclosed	None	None	None	Group signature / HE

● Strong / Supported
● Medium / Partial
○ Weak / Not supported
✓ Yes
✗ No

Figure 17. Qualitative Functional Comparison of Different Data-Sharing Scheme

As shown in Fig. 17, the centralized scheme is operationally simple but provides weak data minimization and auditability; the basic blockchain scheme strengthens tamper evidence but lacks minimum-disclosure verification and collaborative risk detection; and the single-encryption scheme protects stored content but provides only medium access control and weak joint-risk and audit capabilities. The proposed scheme has the broadest functional coverage. Figure 17 is intentionally qualitative. It does not claim numerical superiority across different implementations; the

proposed scheme's measured latency and throughput are reported separately in Fig. 10.

5.3.4.1 Component-Level Overhead and Ablation Interpretation

Table 8 combines the separately measured component costs and functional ablations. The values have different units and should be interpreted as a latency budget rather than summed mechanically, because several operations execute off chain or in parallel.

Table 8. Component-level implementation evidence and overhead

Component removed / measured	Security or functional effect	Measured overhead / evidence
Blockchain evidence and audit	No tamper-evident event history or fast traceability	12 ms hash deposit; 138 ms retrieval at 20,000 records
Attribute-based authorization	Coarse permission control; unauthorized requests may pass	58 ms encryption; 46 ms decryption; 100% interception with policy
Proxy re-encryption	Requires decrypt-and-reencrypt workflow	67 ms per authorization conversion

Zero-knowledge verification	Requires disclosure of full verification material	86 ms end-to-end verification; 26.37 +/- 0.32 ms proof arithmetic
Secure aggregation	Coordinator can inspect individual model updates	0.83 s at 50,000 records; accuracy changes from 91.8% to 91.1%
Smart-contract audit	Manual state transitions and incomplete accountability	100% audit-record generation in 1,000 functional cases
Group-signature interface / homomorphic-encryption interface	Optional anonymity / ciphertext-computation functions are not part of the measured prototype	Conceptual or interface-level only; excluded from all numerical claims

5.3.5 Traceability Efficiency Test

To verify the efficiency of the blockchain audit mechanism in abnormality tracing, the experiment sets 1,000, 5,000, 10,000, and 20,000 audit records, and tests the average

traceability time of traditional log retrieval and the proposed on-chain index retrieval. The specific results are shown in Fig. 18.




Figure 18. Traceability Retrieval Efficiency under Different Audit-Record Scales

As shown in Fig. 18, as the number of audit records increases, the time required for traditional log retrieval grows rapidly. In contrast, the proposed scheme establishes an on-chain audit index through application numbers, event numbers, and hash indexes, enabling faster location of relevant business records. When the number of audit records is 20,000, the average retrieval time of the proposed on-chain index is 138 ms, which is approximately 81.6% lower than that of traditional log retrieval. This indicates that the proposed scheme can effectively improve traceability efficiency in scenarios such as abnormal applications, unauthorized access, and abnormal fund disbursement.

5.4 Experimental Conclusions and Discussion

Functional tests verify access control, minimized disclosure, contract execution, and audit completeness. Component tests quantify implemented privacy services and exclude interface-

level group signatures and homomorphic encryption from numerical claims. Horizontal-FL robustness tests show degradation under regional non-IID partitioning and the vulnerability of standard FedAvg, together with the conditional benefit of clipping and robust aggregation under controlled collusion. Figure 17 is a qualitative functional comparison only; it is not a cross-platform latency or throughput benchmark.

Under the original controlled Fabric configuration, throughput reaches 718 TPS and average latency is 236 ms at 400 concurrent requests; these are point estimates. The end-to-end zero-knowledge proof generation time is 143 ms, while offline commitment precomputation reduces the online arithmetic portion for four 2048-bit subproofs from 8.03 +/- 0.34 ms to 0.06 +/- 0.01 ms. Strong non-IID data retain 82.69 +/- 0.09% F1, and the robust prototype retains approximately 84% F1 under one or two colluding sign-flip clients. These tests strengthen reproducibility but should not be generalized beyond the stated synthetic and prototype settings.

Future work will prioritize a multi-district consortium pilot in which local systems remain systems of record, a provincial or regional governance committee manages contract and certificate policies, and cross-region risk models exchange only approved summaries. Evaluation will include desensitized local records under formal approval, seasonal and policy drift, backdoor and adaptive poisoning, collusion beyond two nodes, disaster recovery, energy and hardware cost, and data-subject lifecycle requests.

6. Conclusion

This paper presents a scenario-specific integration framework for privacy-preserving childcare-subsidy data sharing. Its novelty lies in coordinating established consortium-blockchain, off-chain ciphertext, minimum-disclosure proof, dynamic authorization, federated risk-control, and audit mechanisms around a regulated application-to-disbursement workflow; it does not claim a new cryptographic primitive. The revision grounds the workflow in peer-reviewed childcare-subsidy administration literature, benchmarks the data model against a peer-reviewed administrative-data evaluation framework, distinguishes implemented and conceptual modules, adds a standard pseudocode algorithm and deployment roadmap, and strengthens evaluation with repeated non-IID, proof-precomputation, and collusion tests. The original controlled prototype reports 718 TPS, 236 ms latency, 91.8% federated-learning accuracy, and 138 ms retrieval for 20,000 audit records. The added tests show that offline commitment precomputation can substantially shorten the online proof path and that robust aggregation can mitigate controlled sign-flip attacks. These findings support technical feasibility while confirming that authorized local data, production-grade cryptography, multi-orderer deployment, and institutional governance are prerequisites for real public-sector adoption.

Funding

This work was supported by the Key Discipline Major Project of Liaoning Social Science Planning Fund (Grant No. L24ZD041), entitled “Research on the Long-term Mechanism of Rural Elderly-care Governance Under the Background of Common Prosperity”.

References

- [1] Jenkins J M, Nguyen T. Keeping Kids in Care: Reducing Administrative Burden in State Child Care Development Fund Policy[J]. *Journal of Public Administration Research and Theory*, 2022, 32(1): 23-40. DOI: 10.1093/jopart/nuab020.
- [2] Morrissey T W, Heflin C M, Fannin W C. Room to grow: examining participation and stability in child care subsidies using state administrative data[J]. *Early Childhood Research Quarterly*, 2023, 62: 150-162. DOI: 10.1016/j.ecresq.2022.08.003.
- [3] Horak S, Ward C. Evaluating a state child care assistance program using administrative data[J]. *Evaluation and Program Planning*, 2022, 92: 102094. DOI: 10.1016/j.evalprogplan.2022.102094.
- [4] Casino F, Dasaklis T K, Patsakis C. A systematic literature review of blockchain-based applications: Current status, classification and open issues[J]. *Telematics and Informatics*, 2019, 36: 55-81. DOI: 10.1016/j.tele.2018.11.006.
- [5] Monrat A A, Schelén O, Andersson K. A survey of blockchain from the perspectives of applications, challenges, and opportunities[J]. *IEEE Access*, 2019, 7: 117134-117151. DOI: 10.1109/ACCESS.2019.2936094.
- [6] Zheng Z, Xie S, Dai H N, Chen X, Wang H. Blockchain challenges and opportunities: A survey[J]. *International Journal of Web and Grid Services*, 2018, 14(4): 352-375. DOI: 10.1504/IJWGS.2018.095647.
- [7] Ølnes S, Ubacht J, Janssen M. Blockchain in government: Benefits and implications of distributed ledger technology for information sharing[J]. *Government Information Quarterly*, 2017, 34(3): 355-364. DOI: 10.1016/j.giq.2017.09.007.
- [8] Elisa N, Yang L, Chao F, Cao Y. A framework of blockchain-based secure and privacy-preserving e-government system[J]. *Wireless Networks*, 2023, 29(3): 1005-1015. DOI: 10.1007/s11276-018-1883-0.
- [9] Zhang Y, Kasahara S, Shen Y, Jiang X, Wan J. Smart contract-based access control for the Internet of Things[J]. *IEEE Internet of Things Journal*, 2019, 6(2): 1594-1605. DOI: 10.1109/IJOT.2018.2847705.
- [10] Batubara F R, Ubacht J, Janssen M. Challenges of blockchain technology adoption for e-government: A systematic literature review[C]//*Proceedings of the 19th Annual International Conference on Digital Government Research: Governance in the Data Age*. New York: ACM, 2018: Article 76, 9 pages. DOI: 10.1145/3209281.3209317.
- [11] Tan E, Mahula S, Crompvoets J. Blockchain governance in the public sector: A conceptual framework for public management[J]. *Government Information Quarterly*, 2022, 39(1): 101625. DOI: 10.1016/j.giq.2021.101625.
- [12] Janssen M, Weerakkody V, Ismagilova E, Sivarajah U, Irani Z. A framework for analysing blockchain technology adoption: Integrating institutional, market and technical factors[J]. *International Journal of Information Management*, 2020, 50: 302-309. DOI: 10.1016/j.ijinfomgt.2019.08.012.
- [13] Piao C, Hao Y, Yan J, Jiang X. Privacy preserving in blockchain-based government data sharing: A Service-On-Chain (SOC) approach[J]. *Information Processing & Management*, 2021, 58(5): 102651. DOI: 10.1016/j.ipm.2021.102651.
- [14] Elisa N, Yang L, Chao F, Naik N, Boongoen T. A secure and privacy-preserving e-government framework using blockchain and artificial immunity[J]. *IEEE Access*, 2023, 11: 8773-8789. DOI: 10.1109/ACCESS.2023.3239814.
- [15] Hassija V, Chamola V, Saxena V, Jain D, Goyal P, Sikdar B. A survey on IoT security: Application areas, security threats, and solution architectures[J]. *IEEE Access*, 2019, 7: 82721-82743. DOI: 10.1109/ACCESS.2019.2924045.
- [16] Dwivedi A D, Srivastava G, Dhar S, Singh R. A decentralized privacy-preserving healthcare blockchain for IoT[J]. *Sensors*, 2019, 19(2): 326. DOI: 10.3390/s19020326.
- [17] Hasselgren A, Kralevska K, Gligoroski D, Pedersen S A, Faxvaag A. Blockchain in healthcare and health sciences—A scoping review[J]. *International Journal of Medical Informatics*, 2020, 134: 104040. DOI: 10.1016/j.ijmedinf.2019.104040.
- [18] McMahan B, Moore E, Ramage D, Hampson S, Agüera y Arcas B. Communication-efficient learning of deep networks from decentralized data[C]//*Proceedings of the 20th International Conference on Artificial Intelligence and Statistics. Proceedings of Machine Learning Research*, 2017, 54: 1273-1282.

- [19] Yang Q, Liu Y, Chen T, Tong Y. Federated machine learning: Concept and applications[J]. *ACM Transactions on Intelligent Systems and Technology*, 2019, 10(2): 12:1-12:19. DOI: 10.1145/3298981.
- [20] Kairouz P, McMahan H B, Avent B, Bellet A, Bennis M, Bhagoji A N, Bonawitz K, Charles Z, Cormode G, Cummings R, et al. Advances and open problems in federated learning[J]. *Foundations and Trends in Machine Learning*, 2021, 14(1-2): 1-210. DOI: 10.1561/22000000083.
- [21] Li T, Sahu A K, Talwalkar A, Smith V. Federated learning: Challenges, methods, and future directions[J]. *IEEE Signal Processing Magazine*, 2020, 37(3): 50-60. DOI: 10.1109/MSP.2020.2975749.
- [22] Mothukuri V, Parizi R M, Pouriyeh S, Huang Y, Dehghantanha A, Srivastava G. A survey on security and privacy of federated learning[J]. *Future Generation Computer Systems*, 2021, 115: 619-640. DOI: 10.1016/j.future.2020.10.007.
- [23] Liu K, Yan Z, Liang X, Kantola R, Hu C. A survey on blockchain-enabled federated learning and its prospects with digital twin[J]. *Digital Communications and Networks*, 2024, 10(2): 248-264. DOI: 10.1016/j.dcan.2022.08.001.
- [24] Wu L, Ruan W, Hu J, He Y. A survey on blockchain-based federated learning[J]. *Future Internet*, 2023, 15(12): 400. DOI: 10.3390/fi15120400.
- [25] Hu K, Gong S, Zhang Q, Seng C, Xia M, Jiang S. An overview of implementing security and privacy in federated learning[J]. *Artificial Intelligence Review*, 2024, 57: 204. DOI: 10.1007/s10462-024-10846-8.
- [26] Guembe B, Misra S, Azeta A. Privacy issues, attacks, countermeasures and open problems in federated learning: A survey[J]. *Applied Artificial Intelligence*, 2024, 38(1): 2410504. DOI: 10.1080/08839514.2024.2410504.
- [27] Bonawitz K, Ivanov V, Kreuter B, Marcedone A, McMahan H B, Patel S, Ramage D, Segal A, Seth K. Practical secure aggregation for privacy-preserving machine learning[C]//*Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. New York: ACM, 2017: 1175-1191. DOI: 10.1145/3133956.3133982.
- [28] Guo Z, Sun Y, Zhang X, Wu L. An Efficient Privacy-Preserving Secure Aggregation Scheme for Federated Learning with Input Verification and Dropout Resistance[J]. *EAI Endorsed Transactions on Scalable Information Systems*, 2026, 12(8). DOI: 10.4108/eetsis.11991.
- [29] Li Q, Wen Z, He B. Practical federated gradient boosting decision trees[J]. *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020, 34(4): 4642-4649. DOI: 10.1609/aaai.v34i04.5895.
- [30] Zhang C, Xie Y, Bai H, Yu B, Li W, Gao Y. A survey on federated learning[J]. *Knowledge-Based Systems*, 2021, 216: 106775. DOI: 10.1016/j.knsys.2021.106775.
- [31] Sahai A, Waters B. Fuzzy identity-based encryption[C]//*Advances in Cryptology-EUROCRYPT 2005*. Lecture Notes in Computer Science, vol. 3494. Berlin: Springer, 2005: 457-473. DOI: 10.1007/11426639_27.
- [32] Bethencourt J, Sahai A, Waters B. Ciphertext-policy attribute-based encryption[C]//*Proceedings of the 2007 IEEE Symposium on Security and Privacy*. Oakland: IEEE, 2007: 321-334. DOI: 10.1109/SP.2007.11.
- [33] Lewko A, Waters B. Decentralizing attribute-based encryption[C]//*Advances in Cryptology-EUROCRYPT 2011*. Lecture Notes in Computer Science, vol. 6632. Berlin: Springer, 2011: 568-588. DOI: 10.1007/978-3-642-20465-4_31.
- [34] Ateniese G, Fu K, Green M, Hohenberger S. Improved proxy re-encryption schemes with applications to secure distributed storage[J]. *ACM Transactions on Information and System Security*, 2006, 9(1): 1-30. DOI: 10.1145/1127345.1127346.
- [35] Ben-Sasson E, Chiesa A, Garman C, Green M, Miers I, Tromer E, Virza M. Zerocash: Decentralized anonymous payments from Bitcoin[C]//*Proceedings of the 2014 IEEE Symposium on Security and Privacy*. San Jose: IEEE, 2014: 459-474. DOI: 10.1109/SP.2014.36.
- [36] Goldwasser S, Micali S, Rackoff C. The knowledge complexity of interactive proof systems[J]. *SIAM Journal on Computing*, 1989, 18(1): 186-208. DOI: 10.1137/0218012.
- [37] Gennaro R, Gentry C, Parno B, Raykova M. Quadratic span programs and succinct NIZKs without PCPs[C]//*Advances in Cryptology-EUROCRYPT 2013*. Lecture Notes in Computer Science, vol. 7881. Berlin: Springer, 2013: 626-645. DOI: 10.1007/978-3-642-38348-9_37.
- [38] Paillier P. Public-key cryptosystems based on composite degree residuosity classes[C]//*Advances in Cryptology-EUROCRYPT 1999*. Lecture Notes in Computer Science, vol. 1592. Berlin: Springer, 1999: 223-238. DOI: 10.1007/3-540-48910-X_16.
- [39] Gentry C. Fully homomorphic encryption using ideal lattices[C]//*Proceedings of the 41st Annual ACM Symposium on Theory of Computing*. New York: ACM, 2009: 169-178. DOI: 10.1145/1536414.1536440.
- [40] Boneh D, Boyen X, Shacham H. Short group signatures[C]//*Advances in Cryptology-CRYPTO 2004*. Lecture Notes in Computer Science, vol. 3152. Berlin: Springer, 2004: 41-55. DOI: 10.1007/978-3-540-28628-8_3.
- [41] Song D X, Wagner D, Perrig A. Practical techniques for searches on encrypted data[C]//*Proceedings of the 2000 IEEE Symposium on Security and Privacy*. Berkeley, CA, USA: IEEE Computer Society, 2000: 44-55. DOI: 10.1109/SECPRI.2000.848445.
- [42] Dwork C. Differential privacy[C]//*Proceedings of the 33rd International Colloquium on Automata, Languages and Programming, Part II*. Lecture Notes in Computer Science, vol. 4052. Berlin: Springer, 2006: 1-12. DOI: 10.1007/11787006_1.
- [43] Dwork C, Roth A. The algorithmic foundations of differential privacy[J]. *Foundations and Trends in Theoretical Computer Science*, 2014, 9(3-4): 211-407. DOI: 10.1561/04000000042.
- [44] Chentharu S, Ahmed K, Whittaker F. Privacy-Preserving Data Sharing using Multi-layer Access Control Model in Electronic Health Environment[J]. *EAI Endorsed Transactions on Scalable Information Systems*, 2019, 6(22): e6. DOI: 10.4108/eai.13-7-2018.159356.
- [45] Schnorr C P. Efficient identification and signatures for smart cards[C]//*Advances in Cryptology-CRYPTO '89*. Lecture Notes in Computer Science, vol. 435. Berlin: Springer, 1990: 239-252. DOI: 10.1007/0-387-34805-0_22.