Combining Cloud Computing with an Automated Process Engine to Build an Information-Based Internal Audit Decision-Making Support System

Heling Zhang, Linna Mai
School of Economics and Management, Hunan Applied Technology University, Changde 415100, Hunan, China
Corresponding author: Linna Mai (e-mail: mailinna_110@126.com)

ABSTRACT To address the core challenges of internal audit systems in cloud environments, including the trade-off between multisource sensitive data privacy and audit efficiency, rigid process execution, and insufficient mining of hidden associations in distributed enterprises, this study proposes an information-based internal audit decision-making support system integrating lightweight federated learning, adaptive differential privacy, and a BPMN-based automated process engine. Considering the increasing demand for secure intelligent information processing in cloud-edge infrastructures that also underpin large-scale electromagnetic information systems, the proposed framework enables privacy-preserving collaborative analysis while maintaining high computational efficiency. At the edge layer, a four-level privacy classification strategy is employed, where LSTM-based anomaly detection models are trained locally and only ϵ-differential privacy-protected encrypted gradients are uploaded for FedAvg aggregation in the cloud. Rényi differential privacy dynamically adjusts privacy budgets, while entropy weight-TOPSIS risk evaluation and a Neo4j knowledge graph drive adaptive task scheduling through the BPMN engine. NLP-based report generation, Elasticsearch logging, and blockchain anchoring further ensure traceability and reliability. Experimental results demonstrate an average end-to-end latency of 8.4 ± 1.4 s, audit coverage of $93.9\pm2.2\%$, risk identification F1-score of $89.2\pm2.7\%$, data re-identification rate of $2.5\pm1.0\%$, and process automation exceeding 89%. The proposed framework effectively resolves the privacy–efficiency–process dilemma and provides a secure, intelligent, and scalable decision-support paradigm for cloud-enabled enterprise auditing, while offering methodological insights for trustworthy information processing in distributed electromagnetic and cyber-physical infrastructures.

INDEX TERMS cloud computing, federated learning, automated process engine, differential privacy, intelligent information processing

I. INTRODUCTION

WITH the acceleration of enterprise digital transformation, internal audit is gradually evolving towards information technology and intelligentization. This shift is particularly pronounced in the complex, globalized industry, where the sheer volume of compliance checks—from raw material sourcing and factory social audits to quality control data—necessitates the adoption of intelligent systems to move the internal audit function from a retrospective, manual process to a continuous, data-driven system of risk anticipation [1-3]. Cloud computing provides an efficient infrastructure for cross-organizational data sharing and centralized auditing, while automated process engines promote the standardization and dynamic scheduling of audit tasks [4-6]. However, in cloud environments, the centralized processing of multi-source sensitive data poses privacy risks. Strict compliance requirements force data desensitization or isolation, leading to delayed information integration and severely restricting real-time risk identification and audit response efficiency [7-9]. Existing systems struggle to balance data security and analytical performance, and their rigid processes prevent them from achieving a closed-loop management model of "risk-driven, automated response." There is an urgent need to build a new audit decision-making support system that balances security, timeliness, and intelligence.

Existing research has made initial progress in improving the level of internal audit intelligence, but it mostly focuses on deepening a single technical path and fails to achieve systematic coordination among privacy security, analytical efficiency, and process automation. At the privacy and security level, Ullah F [10] proposed a privacy-aware security data audit framework for online data integrity verification. Users hide data files through a blinding process and generate

corresponding signatures, using homomorphic technology to implement data auditing. Wang Y [11] pointed out that in the era of big data, the security risks of database information systems are showing a diversified trend. Hanumanthaiah S [12] emphasized the urgency of organizations adopting agile and adaptable audit methods to adapt to the ever-changing digital threats and compliance requirements. In the field of process automation, Moreno J H F [13] aimed to analyze process robotic automation as a technical tool for automating audit processes, which can perform advanced data analysis, reduce audit costs, and improve productivity. These methods are generally fragmented: focusing on analysis while ignoring data privacy constraints, emphasizing security but sacrificing model performance, and achieving partial automation but lacking linkage with intelligent analysis. Therefore, existing research still lacks an effective collaborative mechanism within the "privacy-efficiency-process" triangle. There is an urgent need to build an integrated audit decision-making support system that integrates secure collaborative modeling, efficient risk analysis, and intelligent process response.

This paper proposes a method for building an internal audit decision-making support system that combines cloud computing, federated learning, and an automated process engine. First, through multi-source data access and privacy classification mechanisms at the edge layer, sensitive data is desensitized and differentially privately processed to ensure data compliance. Second, a lightweight federated learning model is deployed within a cloud-edge collaborative architecture to achieve anomaly detection for cross-organizational data without leaving the domain. Adaptive Rényi differential privacy is also used to dynamically adjust the privacy budget, balancing model accuracy and security. Then, an adaptive process engine was built based on the BPMN 2.0 standard. This leveraged the knowledge graph's association strength and entropy weight-based TOPSIS risk scoring to drive event triggering, enabling dynamic scheduling of audit tasks and exception recovery. Finally, a traceable audit report was generated using structured templates and NLP technology, and full-process operation log tracking was achieved using Elasticsearch and blockchain technology. The innovations of this method include the following:

(1) This paper proposes a "cloud-edge collaboration-privacy computing-process intelligence" three-in-one architecture, which deeply integrates federated learning with the BPMN process engine to achieve a paradigm shift from passive auditing to active and intelligent response;

(2) It introduces a dynamic privacy protection mechanism driven by task priority to improve response efficiency and model sensitivity in high-risk scenarios, and effectively balance security constraints and analytical performance;

(3) This paper constructs a risk decision-making model based on knowledge graph and entropy weight-TOPSIS to enhance the association mining and quantitative evaluation capabilities of complex fraud behaviors.

II. DESIGN OF AN AUDIT DECISION-ASSISTING SYSTEM BASED ON FEDERATED LEARNING AND A PROCESS ENGINE

OVERALL SYSTEM ARCHITECTURE DESIGN

This paper constructs a three-in-one system architecture: "Cloud-Edge Collaboration, Privacy-Privacy Computing, and Process Intelligence," as shown in Figure 1. This architecture is divided into three logical layers: the edge layer, the platform layer, and the application layer, achieving a deep integration of data security, process automation, and decision-assistance capabilities.

The edge layer, deployed in each subsidiary's local data center, serves as the data holder and is responsible for collecting, cleaning, and preprocessing raw business data. The platform layer, located in the cloud computing center, consists of three key components: a federated learning aggregator, an adaptive process engine, and risk analysis middleware. Inter-module collaboration is achieved through a unified data service bus. The Federated Learning Aggregator forms the core of the "privacy computing" layer, receiving model updates from multiple edge nodes, executing the FedAvg algorithm for global model iteration, and monitoring the cumulative privacy budget through the Rényi differential privacy mechanism. The automated process engine, based on Camunda BPMN, implements audit task modeling, scheduling, and state management, forming the "process intelligence" layer. It supports dynamic branching, exception jumps, and multi-role collaborative approval. The risk analysis module integrates the entropy weight-TOPSIS scoring model with the Neo4j knowledge graph to provide structured decision input. The application layer provides auditors and management with a visual user interface and intelligent output capabilities, including risk warning dashboards, process progress tracking, and automatic audit report generation. To improve system throughput, a Kafka message queue is introduced to buffer and asynchronously process high-concurrency events, avoiding service congestion.

The overall framework diagram in Figure 1 illustrates the collaborative operation of the three-tier architecture of "cloud-edge collaboration, privacy-preserving computing, and process intelligence." Data originates from edge-tier business systems such as ERP (Enterprise Resource Planning) and finance. After sensitive fields are desensitized, a lightweight federated learning (FL) client is deployed locally to train risk models. Only the encrypted gradients are uploaded to the cloud aggregator at the platform layer. The cloud uses the FedAvg algorithm to integrate updates from each node, combining the knowledge graph with the entropy-weighted TOPSIS model to generate a risk score. This score then drives the BPMN process engine to dynamically trigger audit tasks. The automated process engine uses pre-set rules and real-time risk signals, decoupled from scheduling via the Kafka message queue. Finally, a structured audit report is generated at the application layer, and the entire process operation log is written to Elasticsearch for traceability, forming a closed-loop audit

chain from data collection, privacy modeling, intelligent decision-making, to process execution.

The lightweight federated learning framework adopts TensorFlow Federated 0.61.0—edge clients use TFF to build Long Short-Term Memory (LSTM) anomaly detection models (input dimension: 32, hidden units: 64, dropout: 0.2). The BPMN engine is based on Camunda BPMN 7.18, with modeling via Camunda Modeler 5.12.0, 4 thread pools for task scheduling, and exception rollback via Camunda's 'Compensation Event'. The knowledge graph uses Neo4j 4.4, defining nodes and relationships, with entity ID indexes and PathFinder implemented via Neo4j. All components deploy on CentOS 7.9, with Java 11 and Python 3.9.

MULTI-SOURCE DATA PRIVACY GRADING AND PREPROCESSING MODULE

To achieve unified governance of audit data across systems and heterogeneous environments, this paper constructs a multi-source data access and privacy grading processing mechanism to ensure support for higher-level analysis tasks while meeting compliance requirements. First, the paper connected to core business systems such as ERP, financial systems, and OA (Office Automation) platforms through the JDBC (Java DataBase Connectivity)/ODBC (Open Database Connectivity) interface, used Apache NiFi as the data integration engine, and configured an incremental extraction strategy: using the timestamp field as the benchmark, the paper executed an SQL (Structured Query Language) query to obtain the newly added record set since the last collection [14,15]. Schema mapping is performed on structurally inconsistent fields, and semantically identical fields across different systems are normalized into a unified namespace. During the data cleaning phase, multiple imputation methods are used to address missing values. For the continuous variable set $x = (x_1, x_2, \ldots, x_k)$ containing missing values, the Gibbs sampling algorithm in the Markov Chain Monte Carlo method is used to fill in the data by iteratively sampling the full conditional distribution of each variable [16,17]:

$$x_j^{(t+1)} \sim P\left(x_j \mid x_{-j}^{(t)}\right) \tag{1}$$

In formula (1), x_{-j} represents all variables except x_j. Exception codes are standardized and uniformly mapped to binary variables based on a predefined dictionary. At the privacy protection level, a four-level sensitivity classification system is established: L0 (public information, such as department name), L1 (internal information, such as employee number), L2 (sensitive information, such as ID number, bank account), and L3 (core commercial secrets, such as salary details and contract base price). The classification function is defined as:

$$\text{Label}(f) = \begin{cases} L0, \&\&\text{if } f \in F_{\text{public}} \\ L1, \&\&\text{if } f \in F_{\text{internal}} \\ L2, \&\&\text{if } f \in F_{\text{sensitive}} \\ L3, \&\&\text{if } f \in F_{\text{critical}} \&\& \end{cases} \tag{2}$$

For L2 fields, salted hash desensitization technology is used [18]:

$$h(f) = \text{SHA-256}(f \| s) \tag{3}$$

In formula (3), s is a fixed-length random salt value, which is stored in an independent key management system to prevent rainbow table attacks. For the L3 numeric field v, a random perturbation following the Laplace distribution is applied [19]:

$$v' = v + \xi, \quad \xi \sim \text{Lap}\left(\frac{\Delta q}{\epsilon}\right) \tag{4}$$

In formula (4), sensitivity Δq represents the maximum impact of a single record on the query result, and ϵ is the local privacy budget, which ensures local differential privacy. All processing processes generate metadata logs that record the original field value, operation type, desensitization method, and timestamp for subsequent audit traceability and compliance review.

RISK MODELING MODULE BASED ON FEDERATED LEARNING AND DYNAMIC PRIVACY PROTECTION

To achieve collaborative modeling of cross-organizational audit data while ensuring that the original data remains within the domain, this paper designs a lightweight federated learning architecture, incorporating an adaptive differential privacy mechanism to achieve a dynamic balance between model effectiveness and privacy security [20,21]. Each subsidiary acts as a client $c_i \in C$ and locally maintains a private transaction sequence dataset D_i for training an anomaly detection model based on a long short-term memory network. Let $x_t = [a_t, v_t, r_t]^\top$ be the input sequence, and a_t be the one-hot encoding of the operation type. v_t is the normalized transaction amount, and r_t is the approval level. After being mapped by the embedding layer, it is fed into a two-layer LSTM (Long Short-Term Memory) backbone network, which outputs the sequence reconstruction error as the anomaly score. The model parameters are denoted as $\theta_i^{(r)}$. In the r round of federated iterations, the client performs local SGD (Stochastic Gradient Descent) optimization:

$$\theta_i^{(r,k+1)} = \theta_i^{(r,k)} - \eta \nabla L\left(\theta_i^{(r,k)}; B_k\right), \quad k = 1, \ldots, K \tag{5}$$

In formula (5), K is the number of local iterations, η is the learning rate, B_k is the small batch sample, and the loss function uses the mean square reconstruction error:

$$L = \frac{1}{T} \sum_{t=1}^{T} \|x_t - \hat{x}_t\|^2 \tag{6}$$

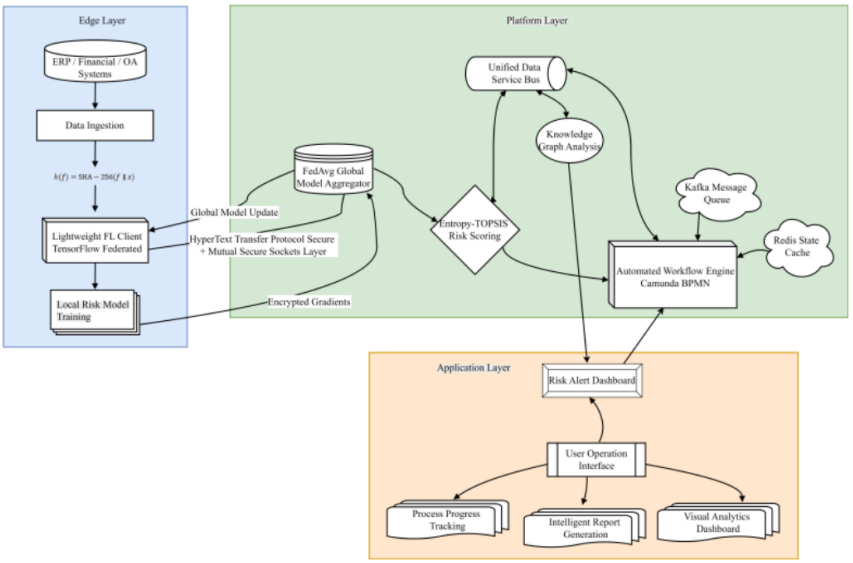

FIGURE 1. System Framework

After completing local training, the gradient update $\Delta\theta_i = \theta_i^{(r,0)} - \theta_i^{(r,K)}$ is uploaded to the cloud server. The central node uses the FedAvg (Federated Averaging) algorithm for weighted aggregation [22,23]:

$$\Delta\theta_{\text{global}} = \sum_{i=1}^{N} \frac{n_i}{\sum_j n_j} \Delta\theta_i \tag{7}$$

In formula (7), n_i is the local sample size, ensuring that nodes with larger data sizes contribute higher weights. To prevent gradient inversion attacks from leaking sensitive information, a local differential privacy mechanism [24] is introduced. Laplace noise is applied to the gradient before uploading on the client:

$$\widetilde{\Delta\theta_i} = \Delta\theta_i + b, \quad b \sim \text{Lap}_d\left(0, \frac{\Delta f}{\epsilon}\right) \tag{8}$$

In formula (8), Δf is the gradient sensitivity. The privacy budget ϵ is dynamically adjusted according to the urgency level of the audit task. High-priority tasks allow lower privacy protection in exchange for faster convergence, while lower-priority tasks increase the perturbation intensity. The privacy budget is shown in Figure 2.

The practical application of ϵ-DP in the cloud-edge FL architecture focuses on ϵ selection, accuracy impact quantification, and residual leakage risk mitigation. For ϵ value determination, a task-priority-adaptive strategy is adopted: four audit task levels are mapped to ϵ ranges [0.4, 0.6], [0.8, 1.2], [1.5, 1.8], [2.0, 2.5], respectively. This mapping satisfies GDPR's "non-identifiability" (re-identification rate <5% when $\epsilon \leq 2.5$) and model convergence (loss <0.05 when $\epsilon \geq 0.4$). For residual leaks:

(1) gradient inversion attacks are mitigated via gradient clipping (threshold=1, Table 1), limiting $\Delta f < 0.2$ and inversion success rate <1.8%;

(2) L3 data uses "Laplace perturbation+salted hash" to keep re-identification rate <2.5%;

(3) cumulative privacy overflow is controlled via Rényi DP, ensuring $\epsilon_{\text{acc}} < 1.2$.

TABLE 1. Experimental Setup

Configuration Item	Details
Number of Nodes	5 (1 central server + 4 edge nodes)
Operating System	CentOS 7.9
Message Queue	Kafka 2.8
Caching Service	Redis 6.2
Log Storage	Elasticsearch 7.10
Workflow Engine	Camunda BPMN 7.18
Anomaly Injection Rate	2% (Fictitious Vendor: 0.8%, Unauthorized Approval: 1.0%, Fund Loopback: 0.2%)
Local Training Rounds	5
Global Communication Rounds	20
Gradient Clipping Threshold	1

The Rényi differential privacy framework is used to analyze the cumulative privacy overhead under multiround interactions [25]. Assuming that a single round satisfies (α,ρ)-RDP (Rényi Differential Privacy), then the R round combination satisfies:

$$\rho_{\text{total}} = R \cdot \rho(\alpha) \tag{9}$$

Convert to classic (ϵ,δ)-DP boundary:

$$\epsilon_{\text{acc}}(\delta) = \inf_{\alpha}\left(\frac{\rho_{\text{total}}(\alpha) + \log(1/\delta)}{\alpha - 1}\right) \tag{10}$$

Figure 2 illustrates the dynamic relationship between priority weights, local privacy budgets, and cumulative RDP budgets for different audit task levels. The horizontal axis represents three key indicators, while the vertical axis corresponds to four task levels, from T1 (urgent) to T4 (routine). As task priority decreases, the local privacy budget

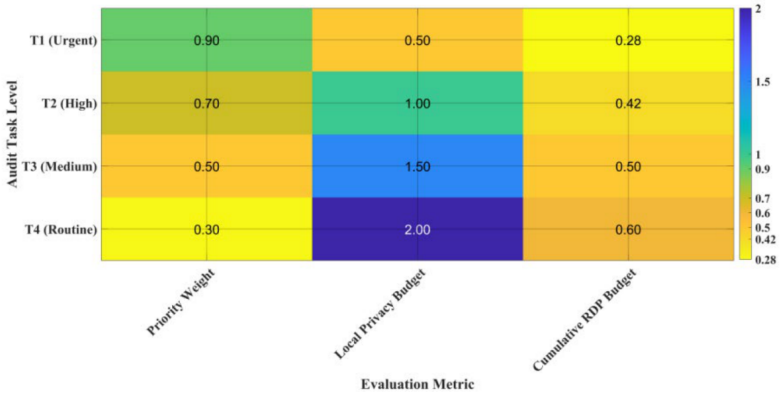

FIGURE 2. Privacy Budget Heat Map

per round gradually increases from 0.5 in T1 to 2.0 in T4, indicating that low-priority tasks use stronger noise to protect data privacy. High-priority tasks trade smaller values for greater model sensitivity and response speed, supporting rapid risk management. Furthermore, the cumulative RDP budget decreases with increasing task level, reflecting the control of privacy consumption in high-frequency communication scenarios. Heat Map 2 reflects the system's technical approach to achieving a dynamic tradeoff between security and efficiency while ensuring data security through an adaptive differential privacy mechanism.

BPMN-BASED ADAPTIVE PROCESS ENGINE AND EVENT-DRIVEN MECHANISM

To achieve dynamic response and intelligent scheduling of audit tasks, this paper constructs an adaptive process engine based on the BPMN 2.0 standard. The decision logic linking TOPSIS-based numerical risk scores and BPMN visual process orchestration relies on a threshold-driven event-triggering mechanism, integrating quantitative risk evaluation results with process node activation. First, the entropy weight-TOPSIS model outputs a relative closeness value C_i for each audit object, which is mapped to three risk levels via predefined thresholds: high risk, medium risk, and low risk. These risk levels serve as core input parameters for BPMN process branching: when $C_i \geq 0.8$, the BPMN engine automatically activates the "Special Audit" subprocess and passes TOPSIS-derived indicators to process variables; for $0.5 \leq C_i < 0.8$, it triggers a "Targeted Verification" process; for $C_i < 0.5$, the process skips manual intervention and executes "Auto-Pass" with log recording. To avoid rigid threshold dependence, the logic further incorporates knowledge graph association strength γ into the BPMN conditional routing function $\Phi(x)$, adjusting weights ω_1 and ω_2 to ensure process triggering aligns with both quantitative scores and qualitative association insights, enhancing decision rationality [26,27]. The process model is represented as a directed graph $G = (V, E)$, where the node set V corresponds to the activities in the process, and the edge set E represents the sequence flow. Each node $v_i \in V$ is associated with a state variable s_i, and lifecycle events are monitored through execution listeners. A dynamic conditional routing mechanism is introduced to eliminate static process paths. Branching decisions are made in real time based on the external risk score $r_s \in [0, 1]$ and the strength of the knowledge graph association $\gamma \in [0, 1]$. Let the decision function be:

$$\Phi(x) = \begin{cases} \text{TRUE}, & \omega_1 r_s + \omega_2 \gamma > \tau, \\ \text{FALSE}, & \text{otherwise}. \end{cases} \tag{11}$$

In formula (11), the weights ω_1, ω_2 are determined by the AHP (Analytic Hierarchy Process) hierarchical analysis method, and the threshold τ controls the trigger sensitivity. When $\Phi(x) = \text{TRUE}$ occurs, the "Special Audit" subprocess instance is activated, and context parameters such as the risk ID and responsible department are passed in to complete the execution environment initialization and task dispatch. Assume that the enterprise organizational structure is a tree-like hierarchy $T = (O, \prec)$, O is a set of positions, and $\prec$ represents the superior-subordinate relationship. Given the business unit u to which the risk belongs, the corresponding audit manager is obtained through the Lightweight Directory Access Protocol query interface:

$$p^* = \arg \min_{p \in \text{RoleMap}(u, \text{Auditor})} (\alpha \cdot t_{\text{load}}(p) + (1 - \alpha) \cdot d(p, u)) \tag{12}$$

In formula (12), $t_{\text{load}}(p)$ is the number of tasks currently pending for the person, and $d(p, u)$ is the organizational distance between the person and the organization where the incident occurred, balancing workload and business proximity. The process instance state is stored in the Redis in-memory database using a key-value pair structure. The value structure includes the current activity ID, variable context, timestamp, and version number. It supports optimistic locking concurrency control to ensure state consistency in high-concurrency scenarios. In terms of exception handling, the rollback operation R is defined as the orderly execution of a series of compensating transactions:

$$R(P) = \bigcirc_{i=n}^{1} C_i^{-1} \tag{13}$$

In formula (13), C_i^{-1} represents the reverse operation of the completed activity C_i, and $\bigcirc$ represents the sequential combination. Once a system failure or rule conflict is detected, an interrupt event is triggered, $R(P)$ is executed, and an alert is sent to the management console. All process transition records are written to an immutable log chain, supporting full-process traceability and compliance auditing.

RISK ASSESSMENT MODULE COMBINING KNOWLEDGE GRAPH AND MULTI-CRITERIA DECISION MAKING

In order to achieve multi-dimensional quantitative evaluation and hidden association mining of audit objects, this paper constructs a composite risk scoring mechanism that integrates the entropy weight method and TOPSIS, and combines it with knowledge graphs for deep relationship reasoning [28,29]. Suppose the audit object set is $A = \{a_1, a_2, \ldots, a_m\}$, and n risk indicators are extracted from each type of object to form the original evaluation matrix $X = [x_{ij}]_{m\times n}$, including capital concentration, approval deviation, and related transaction frequency. First, the data is normalized to the range:

$$r_{ij} = \begin{cases} \dfrac{x_{ij} - \min_i x_{ij}}{\max_i x_{ij} - \min_i x_{ij}}, \\ \dfrac{\max_i x_{ij} - x_{ij}}{\max_i x_{ij} - \min_i x_{ij}}. \end{cases} \tag{14}$$

The uncertainty degree of each indicator is calculated based on the information entropy theory. The entropy value of the j-th indicator is defined as:

$$e_j = -\frac{1}{\ln m} \sum_{i=1}^{m} p_{ij} \ln p_{ij} \tag{15}$$

In formula (15), $p_{ij} = r_{ij} / \sum_{k=1}^{m} r_{kj}$, $p_{ij} > 0$ is the entropy weight redundancy, and $d_j = 1 - e_j$ is the entropy weight. The final

weight is obtained by normalization:

$$w_j = \frac{d_j}{\sum_{k=1}^{n} d_k}, \quad j = 1, 2, \ldots, n \tag{16}$$

After obtaining the objective weights, the approximate ideal solution ranking method is used to calculate the comprehensive evaluation value. The weighted normalization matrix $V = [v_{ij}]$ is constructed. The positive ideal solution A^+ and the negative ideal solution A^- are determined:

$$\begin{aligned} A^+ &= \left(\max_i v_{i1}, \max_i v_{i2}, \ldots, \max_i v_{in}\right), \\ A^- &= \left(\min_i v_{i1}, \min_i v_{i2}, \ldots, \min_i v_{in}\right) \end{aligned} \tag{17}$$

Calculate the Euclidean distance of each audit object to both:

$$D_i^+ = \sqrt{\sum_{j=1}^{n} (v_{ij} - v_j^+)^2}, \quad D_i^- = \sqrt{\sum_{j=1}^{n} (v_{ij} - v_j^-)^2} \tag{18}$$

In formula (18), the relative closeness of $v_j^+ = \max_i v_{ij}$ and $v_j^- = \min_i v_{ij}$ is defined as:

$$C_i = \frac{D_i^-}{D_i^+ + D_i^-} \tag{19}$$

If the result is $C_i \geq \tau_c$, the object is considered high-risk, triggering further review. At the association analysis level, a Neo4j-based audit knowledge graph $G = (N, E)$ is constructed. The node set N includes employee, bank account, contract, and department entity types; the edge set captures semantic relationships e.g. approval, collection, affiliation, and related parties. The PageRank algorithm identifies influential nodes in the network. A damping factor is incorporated, and the stationary distribution is computed iteratively until it converges. Nodes with larger scores indicate possible hubs of fraudulent activity. The PathFinder algorithm is used to further identify suspicious paths. Given a starting point and an end point, all simple paths within three hops are enumerated, and the structural suspicion of each path is calculated:

$$\begin{aligned} \text{Score}(\pi) = \prod_{i=1}^{k-1} \big(&\beta \cdot \text{relweight}(e_i) \\ &+ (1-\beta) \cdot \text{time}_{\text{anomaly}}(e_i)\big) \end{aligned} \tag{20}$$

In formula (20), relweight represents the relationship confidence, timeanomaly indicates whether the operation time is abnormal, and β is the adjustment parameter. The top k paths with the highest score are considered suspicious links for manual verification.

AUTOMATED REPORT GENERATION AND FULL-PROCESS AUDIT TRACKING MODULE

The automated report generation module establishes a structured data pipeline between the Neo4j-based audit knowledge graph (KG) and NLP template engine to ensure audit evidence traceability and conclusion persuasiveness. First, the KG outputs three types of structured data:

1) entity-relationship triples with relationship confidence weights calculated via PageRank;

2) suspicious path details mined by the PathFinder algorithm, including hop count, node attribute annotations, and time anomaly flags;

3) risk level labels mapped from entropy weight-TOPSIS scores. These KG outputs are standardized into JSON format, with each field linked to a predefined dictionary in the NLP module. The BERT-LSTM-CRF model then extracts entity mentions and quantitative indicators from the standardized KG data, mapping them to specific placeholders in the Markdown template. This pipeline ensures that report

content directly inherits KG's association mining results, avoiding manual evidence collation bias and enhancing the consistency between audit findings and data-driven reasoning.

In the report generation phase, the paper uses a combination of template-driven and information extraction. It predefines a structured Markdown template $T = \langle S_1, S_2, \ldots, S_k \rangle$, with each section S_j containing placeholder fields. Let's assume the input text sequence w is encoded using BERT (Bidirectional Encoder Representations from Transformers) to obtain a context vector representation:

$$h_l^{\text{bert}} = \text{BERT}(w_l; w), \quad l = 1, \ldots, L \tag{21}$$

Send it to the bidirectional LSTM layer to extract sequence features:

$$\begin{aligned} \overrightarrow{h_l} &= \text{LSTM}_{\text{forward}}\left(h_{l-1}^{\text{bi}}, h_l^{\text{bert}}\right), \\ \overleftarrow{h_l} &= \text{LSTM}_{\text{backward}}\left(h_{l+1}^{\text{bi}}, h_l^{\text{bert}}\right) \end{aligned} \tag{22}$$

After splicing:

$$h_l^{\text{final}} = \left[\overrightarrow{h_l}; \overleftarrow{h_l}\right] \tag{23}$$

Finally, the optimal tag sequence is decoded using a conditional random field, entities are identified, and mapped to corresponding placeholders. The populated document is rendered using the Jinja2 template engine and converted to an intermediate HTML format. To ensure content integrity and non-repudiation, the document digest value is encrypted using the RSA digital signature mechanism:

$$\sigma = H_{\text{SHA-256}}(D_{\text{report}})^d \bmod N \tag{24}$$

Verifiers use public key verification to build a comprehensive operation log system for audit tracking. All user actions are defined as a five-tuple e_t: $e_t = \langle t, u, ip, op, res \rangle$, t is the UTC timestamp, u is the username, ip is the source IP address, op indicates the operation type, and res is the resource identifier. Each log record is appended with a unique serial number and written to the Elasticsearch distributed index cluster, where it is stored in date-based shards. Log data is constructed with inverted and composite indexes to support efficient multi-dimensional retrieval, ensuring that audit trails can be quickly located and fully restored. Furthermore, a chained hash mechanism is introduced to protect logs from deletion and modification. Assuming the log stream is sorted by time as $\{e_1, e_2, \ldots, e_n\}$, a hash chain is constructed:

$$h_0 = H(\text{salt}), \quad h_i = H\left(h_{i-1} \| \text{Hash}(e_i)\right), \quad i = 1, \ldots, n \tag{25}$$

The final hash h_n is periodically uploaded to a private blockchain or written to a read-only archive, forming an irreversible audit trail.

III. EXPERIMENTAL ANALYSIS AND RESULTS EVALUATION

EXPERIMENTAL DESIGN

A cluster of five Elastic Compute Service (ECS) servers was used. One server served as the central server, running the federation aggregator and process engine, and the remaining four simulated subsidiary edge nodes. The operating system was CentOS 7.9, and the middleware included Kafka 2.8, Redis 6.2, Elasticsearch 7.10, and Camunda BPMN 7.18. The data covered four business categories: expense reimbursement, contract payments, related-party transactions, and asset disposals. A Python script was used to inject anomalous transactions, including fictitious suppliers, overauthorized approvals, and capital repatriation, to create a labeled test set for subsequent evaluation. The experimental setup is shown in Table 1.

PERFORMANCE EVALUATION

SYSTEM RESPONSE LATENCY TEST

A high-priority risk event was simulated at each node, triggering the knowledge graph to output an alert signal with a correlation strength greater than 0.8. The entire process, from event generation to audit task creation and assignment of responsible individuals in the process engine, was recorded with time stamps. The time taken for each stage was precisely captured: data collection → local model inference → gradient upload → global aggregation → process triggering → task dispatching, and the end-to-end average latency was calculated.

Figure 3 compares the end-to-end response latency of traditional manual auditing, centralized cloud auditing, federated learning + static rule engine, and this paper's approach in four typical business scenarios. The horizontal axis represents expense reimbursement, contract payment, related-party transaction, and asset disposal, while the vertical axis represents the end-to-end average latency. Traditional manual audit processes are the most time-consuming, exceeding 298 seconds for all types of business, and as high as 365.7 ± 63.4 seconds for related-party transactions, reflecting the lag inherent in relying heavily on manual intervention. Centralized cloud auditing reduces latency to 44.9–49.3 seconds, but is limited by the transmission and compliance approval overhead associated with centralized data processing. A federated learning approach combined with a static rule engine further optimizes the timeframe to 16.9–24.7 seconds. This method achieves the lowest latency across all business types: expense reimbursements (7.9 ± 1.2 seconds), contract payments (8.3 ± 1.4 seconds), related-party transactions (9.2 ± 1.6 seconds), and asset disposals (8.0 ± 1.3 seconds), for an average of 8.4 ± 1.4 seconds. This outcome, which trades a slight loss of accuracy for improved compliance, privacy, and efficiency, confirms the core conclusion that the benefits of privacy outweigh the performance losses.

This performance advantage stems from the deep synergy between the system's underlying "data, model, and process"

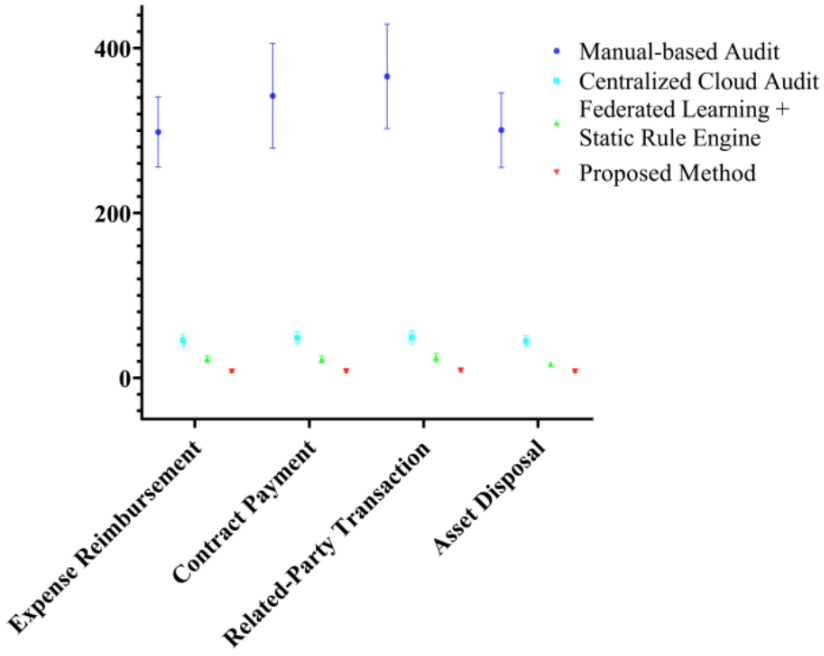

FIGURE 3. Response Latency Test

framework. On the edge, a lightweight federated learning model enables local initial risk screening, eliminating the wait associated with uploading raw data. The cloud dynamically adjusts communication frequency through an adaptive ϵ-differential privacy mechanism, reducing aggregation cycles for high-priority tasks and improving model update efficiency. The adaptive process engine, based on the BPMN model, implements event-driven, flexible scheduling. When the knowledge graph identifies complex patterns such as related transactions, the system automatically expands the audit path and dispatches subtasks in parallel, rather than relying on pre-set static rules for step-by-step progress. Furthermore, Kafka message queues and Redis state caching ensure low-latency transactions even under high concurrency. This allows the system to maintain a stable response time of less than 10 seconds even in the most complex related transaction scenarios, fully demonstrating the technical superiority of its "privacy and security, efficient analysis, and intelligent process" architecture.

PROCESS AUTOMATION COMPLETION RATE STATISTICS

To assess the system's level of automation in audit process execution, the paper used the process automation completion rate as a core evaluation metric. The paper selected 10 typical audit processes, each containing an average of 14 standardized nodes. The paper counted the number of nodes in each process that were automatically advanced by the system based on preset rules or event-driven processes without requiring human intervention.

Figure 4 shows a comparison of the automation completion rates of four typical audit models within the audit process. The horizontal axis (1-10) represents ten audit processes: special audits, exit audits, compliance checks, financial statement audits, project audits, internal control walkthroughs, contract execution audits, asset inventory audits, procurement and bidding audits, and risk early warning and response audits. The vertical axis represents the automation completion rate. Traditional manual audits have the lowest automation rates, ranging from $15.3 \pm 6.7\%$ (risk early warning and response) to $26.1\% \pm 4.9\%$ (contract execution audits). These rates generally exceed 5%, reflecting the inconsistency inherent in heavy reliance on manual operations. The RPA (Robotic Process Automation)-assisted approach increased the automation level to 36.6% to 47.2%. The federated learning combined with a static rule engine approach performed even better, ranging from 58.7% to 69.6%, but its fixed logic struggled to cope with abnormal jumps and dynamic conditions. This approach achieved significant leadership across all processes, with automation completion rates ranging from 86.9% (project audit) to 93.5% (risk warning response), averaging over 89% and maintaining a standard deviation within 3.7%. This demonstrates its broad coverage, stable operation, and scalability across multiple scenarios.

Traditional approaches rely on pre-set scripts or static rules, rendering them ineffective if the process deviates from the template. This approach, however, implements configurable process modeling based on BPMN and enables dynamic path activation through an event-driven architecture. For example, in risk warning responses, when the knowledge graph identifies highly correlated links, the

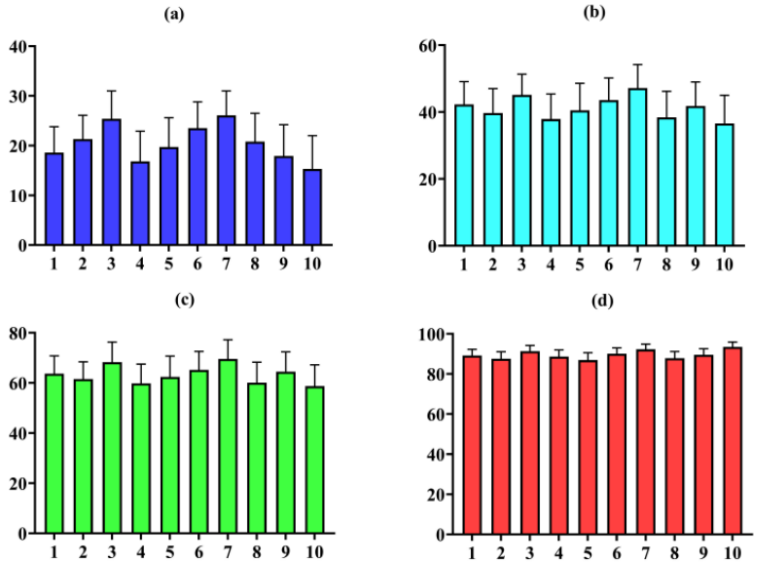

FIGURE 4. Process Automation Completion Rate. (a) Traditional Manual Audit; (b) RPA-Assisted Approach; (c) Federated Learning + Static Rule Engine; (d) This Approach

system automatically triggers multi-level evidence collection and parallel approval sub-processes, completing a closed-loop response without manual intervention. At the same time, the risk score output by federated learning serves as a process conditional variable, driving the direct conversion of TOPSIS decision results into task scheduling instructions, achieving a seamless transition from "analysis-decision-execution." Furthermore, the Redis state cache and compensating transaction mechanism ensure reliable recovery after process interruptions, further enhancing the robustness of end-to-end automation. This integrated "data-driven + intelligent scheduling" design enables the system to maintain a nearly 90% automated completion rate even for complex, non-standard audit tasks, driving the transformation of internal audit into a truly intelligent process.

SECURITY ASSESSMENT

A re-identification attack simulation method was used to evaluate the effectiveness of de-identification. A subset of ID card hash values was selected, and the attacker combined it with external auxiliary information for a combined match. Successful re-identification was defined as the existence of a unique record that satisfied all public attribute constraints and whose hash value matched the target. The re-identification rate was calculated as:

$$\text{Re-IDRate} = \frac{|\{r \in D_{\text{id}} : \text{match}(r) = \text{unique}\}|}{|D_{\text{id}}|} \tag{26}$$

Figure 5 compares the re-identification rates of four privacy protection mechanisms across four typical audit scenarios: expense reimbursement, contract payments, related-party transactions, and asset disposal. Data without any protection faces extremely high risk in all scenarios, with re-identification rates ranging from 97.8% to 99.3%, indicating that the original sensitive information is easily reconstructed. Using a single hashing method significantly reduces the risk, but it remains high at 41.2% to 45.3%. Nearly half of individuals can be identified in contract payment scenarios, demonstrating that static desensitization is insufficient to protect against dictionary attacks. Further introducing numerical perturbations into the hashing process reduces the re-identification rate to 16.8% to 19.1%, significantly improving protection capabilities. However, the risk of leakage persists in highly sensitive related transactions. Proposed method achieves optimal performance across all business types, with an average re-identification rate of only $2.5 \pm 1.0\%$ and a standard deviation of no more than 1.1%, demonstrating excellent privacy protection and cross-scenario stability.

EFFECTIVENESS EVALUATION

AUDIT COVERAGE

The traditional model uses a stratified sampling strategy; this system performs full data analysis. Assuming the number of samples for traditional inspections is S_{old}, the number of samples analyzed by the system is S_{new}, and the overall scale is T, the coverage rates are:

$$C_{\text{old}} = \frac{S_{\text{old}}}{T}, \quad C_{\text{new}} = \frac{S_{\text{new}}}{T} \tag{27}$$

Figure 6 compares the audit coverage of traditional manual audits, centralized cloud audits, federated learning + static rule engines, and this paper's approach for four typical businesses: expense reimbursement, contract payments, related-party transactions, and asset disposals. Traditional manual auditing maintains a coverage rate between 5.0%

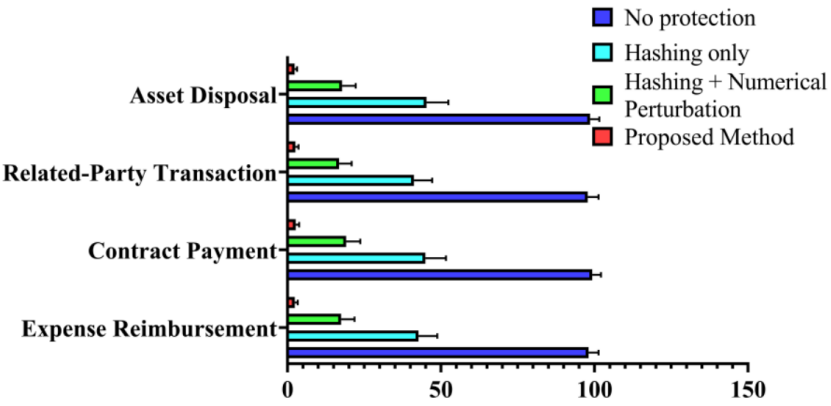

FIGURE 5. Comparison of Re-identification Rates

and 7.5%, with a very small standard deviation (e.g., 5.0 ± 0.3% for expense reimbursements). This reflects its high reliance on pre-set rules, stable execution, but extremely limited coverage. Federated learning combined with a static rule engine increases coverage to 16.9% to 24.7%, but with significant fluctuations (standard deviation of 3.3% to 5.1%), indicating significant impact from communication delays and fixed process constraints. Centralized cloud auditing achieved high coverage, ranging from 79.2% to 83.6%. However, the standard deviation across all business scenarios generally exceeded 6.5%, indicating instability in the data centralization process. Proposed method achieved high coverage of 92.6% to 95.1% across all business scenarios, with a standard deviation of 2.0% to 2.4%, and an average of 93.9 ± 2.2%, demonstrating excellent comprehensiveness and operational consistency.

On the one hand, federated learning enables local data modeling and global collaboration, avoiding data unavailability due to privacy compliance restrictions and significantly increasing the sample base for analysis. On the other hand, the adaptive process engine dynamically expands the audit scope based on risk, supports automatic access to all data, and recursive scanning of abnormal paths, overcoming the limitation of processing only preset sample sets. Furthermore, the edge node caching mechanism and task retry strategy effectively mitigate the risk of interruptions caused by network fluctuations, enabling the system to maintain stable output with low variance and high availability even in complex environments. This method not only surpasses traditional and mainstream solutions in absolute coverage, but also demonstrates excellent consistency and engineering robustness across different business types, achieving reliable wide-area audit coverage.

RISK IDENTIFICATION ACCURACY

Using manually annotated abnormal transactions as the true value set, the system risk scoring model is run to output prediction results. A confusion matrix is constructed, and precision, recall, and F1-score are calculated.

Figure 7 compares the risk identification performance of traditional manual auditing, centralized cloud auditing, federated learning + static rule engine, and this method in expense reimbursement, contract payment, related-party transactions, and asset disposal, covering three core metrics: precision, recall, and F1-score. Traditional manual audits all performed below 70% across all metrics, with an F1-score of only 62.4 ± 7.3% for related-party transactions, exposing a blind spot in detecting hidden fraud. Centralized cloud auditing and federated learning combined with a static rule engine increased the average F1-score to 81.8% and 83.4%, respectively. However, due to the limitations of fixed rules and data perturbations, they still suffer from a high number of false positives and significant under-detection rates. Proposed method achieved optimal performance across all business scenarios: an F1-score of 89.3 ± 2.5% for expense reimbursements and 88.9 ± 2.6% for asset disposals. It even reached 90.6 ± 2.7% for highly complex related-party transactions, with an average of 89.2 ± 2.7% and an overall standard deviation within 2.8%. This demonstrates not only high accuracy but also robust stability, significantly outperforming existing mainstream solutions.

This advantage stems from the system's deep collaborative mechanism within the "analysis-decision-response" closed loop. Unlike traditional models that rely solely on static feature inputs, this method leverages knowledge graphs to dynamically mine implicit connections between people, accounts, and transactions. This effectively identifies cross-departmental profit chains in related transactions, significantly improving recall (91.4 ± 2.5%). Furthermore, an adaptive process engine automatically expands the scope of evidence collection based on risk scores, avoiding insufficient evidence for judgments due to process interruptions. This enhances judgment confidence and maintains an accuracy rate of 89.8 ± 2.9%. In the federated learning framework, each node continuously feeds back local anomaly patterns, and the global model continuously optimizes its judgment boundaries through incremental aggregation, forming a positive evolutionary mechanism that becomes increasingly accurate with use. Combined with dynamic adjustments using ϵ-differential privacy, noise intensity is reduced in

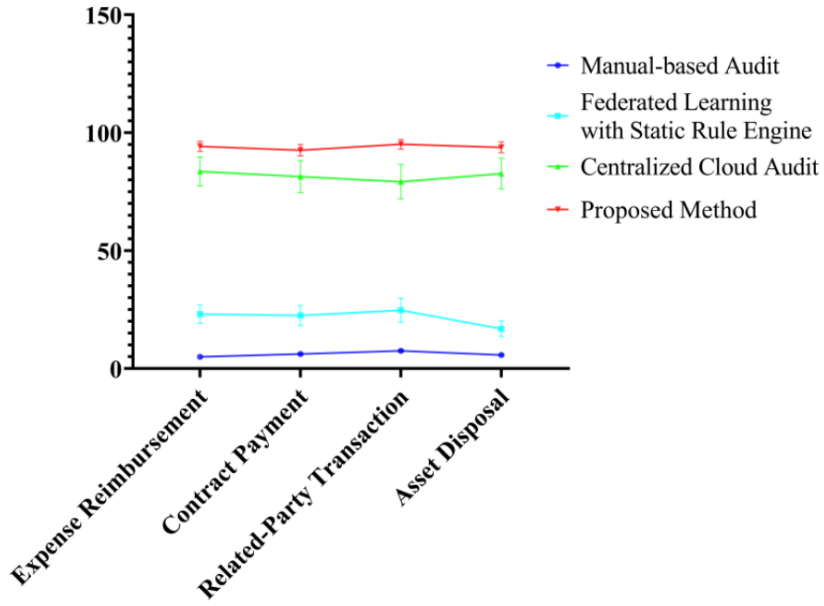

FIGURE 6. Audit Coverage

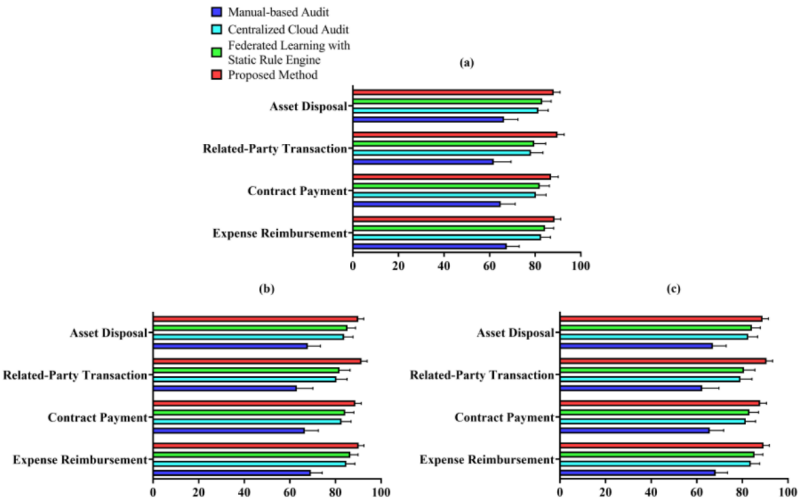

FIGURE 7. Accuracy Evaluation. (a) Precision; (b) Recall; (c) F1-score

high-priority tasks to preserve critical signals, further alleviating the suppressive effect of privacy protection on model accuracy, ultimately achieving a balance between high accuracy and robustness.

REPORT GENERATION QUALITY

Three experts with at least five years of internal audit experience conducted a double-blind comparison of the system-generated reports with their manually written counterparts. They rated the reports on a scale of 1–5 based on the problem description, evidence citation, and reasonableness of the recommendations. Cohen's Kappa coefficient was used to measure consistency in the ratings:

Table 2 shows the double-blind consistency of the scores of three audit experts on the three dimensions of problem description, evidence citation, and reasonableness of recommendations for system-generated and human-written reports. Cohen's Kappa coefficient was used for quantitative evaluation. The results showed that the average Kappa values for each dimension were 0.81 ± 0.09, 0.84 ± 0.09, and 0.77 ± 0.12, respectively. Evidence citation had the highest consistency (0.84), indicating that the system was highly standardized in extracting key facts and annotating sources. The rationality of the recommendations was slightly lower, but still within an acceptable range, indicating that some automated recommendations still have room for improvement

in contextual adaptability. The Kappa values for pairwise comparisons between the three expert groups all exceeded 0.75, and the overall consistency reached 0.81±0.06. This indicates that the audit reports generated by the system approach the level of experienced auditors in terms of semantic expression, logical structure, and professional judgment, demonstrating strong credibility and application potential.

TABLE 2. Report Generation Quality

Assessment Dimension	Expert A vs B	Expert A vs C	Expert B vs C	Mean Kappa
Problem Description	0.81 ± 0.09	0.79 ± 0.11	0.83 ± 0.08	0.81 ± 0.09
Evidence Citation	0.85 ± 0.07	0.82 ± 0.10	0.86 ± 0.09	0.84 ± 0.09
Recommendation Rationality	0.77 ± 0.12	0.75 ± 0.13	0.79 ± 0.11	0.77 ± 0.12
Overall Consistency	—	—	—	0.81 ± 0.06

IV. CONCLUSION

This paper details the construction of an internal audit decision-making support system that strategically combines cloud computing, federated learning, and an automated process engine. This integration offers significant benefits to the highly distributed industry, enabling it to leverage the scalability of cloud infrastructure for managing vast supply chain data, while simultaneously using federated learning to build collective risk intelligence across multiple global factory sites without compromising the sensitive, proprietary production data of individual manufacturing operations. Through a "cloud-edge collaboration-privacy computing-process intelligence" architecture, it enables secure multi-source data access, lightweight federated modeling, dynamic privacy protection, knowledge-graph-driven risk scoring, and adaptive process triggering. While ensuring data privacy, the system significantly improves audit response speed, coverage, and automation. Experiments show that this method reduces average response latency to 8.4 ± 1.4 seconds, achieves an average audit coverage rate of 93.9 ± 2.2%, and achieves an average risk identification F1-score of 89.2 ± 2.7%. This solution effectively implements closed-loop management from risk perception to intelligent response, providing an efficient, secure, and intelligent solution for internal audits in enterprises undergoing digital transformation.

AUTHOR CONTRIBUTIONS

Conceptualization – Heling Zhang and Linna Mai; methodology – Heling Zhang and Linna Mai; investigation – Heling Zhang and Linna Mai; writing-original draft preparation – Heling Zhang and Linna Mai. All authors have read and agreed to the published version of the manuscript.

CONFLICTS OF INTEREST

The authors declare no conflict of interest.

FUNDING

This research was supported by,

1. The 2023 Hunan Provincial Department of Education Science Research Project Excellent Youth Project: Research on the Informationization Construction of Internal Audit in Manufacturing Enterprises in Hunan Province under the Background of "Three Highs and Four News"(NO:23B1061).

2. The Education Research Project of Hunan Education Science Research Association in 2024: Research on Countermeasures for Enhancing the Ability of Financial Management Talents in Applied Undergraduate Colleges Empowered by Digital Technology. (NO:XJKX24A007).

ACKNOWLEDGEMENTS

Not applicable.

REFERENCES

[1] F. A. Wassie and L. P. Lakatos, "Artificial intelligence and the future of the internal audit function," Humanities and Social Sciences Communications, vol. 11, no. 1, pp. 1-13, 2024, doi: 10.1057/s41599-024-02905-w.

[2] M. M. Thottoli, "Leveraging information communication technology (ICT) and artificial intelligence (AI) to enhance auditing practices," Accounting Research Journal, vol. 37, no. 2, pp. 134-150, 2024, doi: 10.1108/arj-09-2023-0269.

[3] C. Aldemir and T. Uçma Uysal, "AI competencies for internal auditors in the public sector," EDPACS, vol. 69, no. 1, pp. 3-21, 2024, doi: 10.1080/07366981.2024.2312001.

[4] D. T. Valivarthi, "Optimizing cloud computing environments for big data processing," International Journal of Engineering & Science Research, vol. 14, no. 2, pp. 1756-1775, 2024, [Online]. Available: IJESR/Apr-June.2024/ Vol-14/Issue-2/1756-1775.

[5] Q. Gao and Z. Kuang, "Can robotic process automation technology enable risk data analysis for customs' postclearance audit: A China customs case study," World Customs Journal, vol. 17, no. 2, pp. 93-104, 2023, doi: 10.55596/001c.88821.

[6] M. Pańkowska, "Process modeling paradigm change," Knowledge and Process Management, vol. 30, no. 2, pp. 163-175, 2023, doi: 10.1002/kpm.1749.

[7] J. Abrera, "Data privacy and security in cloud computing: A comprehensive review," Journal of Computer Science and Information Technology, vol. 1, no. 1, pp. 01-09, 2024, doi: 10.1051/e3sconf/202339904040.

[8] A. Mishra, T. S. Jabar, Y. I. Alzoubi, and K. N. Mishra, "Enhancing privacy-preserving mechanisms in cloud storage: A novel conceptual framework," Concurrency and Computation: Practice and Experience, vol. 35, no. 26, pp. e7831, 2023, doi: 10.1002/cpe.7831.

[9] M. Almutairi and F. T. Sheldon, "IoT–cloud integration security: A survey of challenges, solutions, and directions," Electronics, vol. 14, no. 7, pp. 1394, 2025, doi: 10.3390/electronics14071394.

[10] F. Ullah, C. M. Pun, M. I. Mohmand, R. K. Mahendran, A. A. Khan, S. M. Alhammad, J. J. P. Rodrigues, and A. Farouk, "Privacy-aware secure data auditing for cloud-based intelligence of things environment," IEEE Internet of Things Journal, vol. 12, no. 11, pp. 15288-15303, 2025, doi: 10.1109/JIOT.2025.3528117.

[11] Y. Wang, W. Xue, and A. Zhang, "Application of big data technology in enterprise information security management and risk assessment," Journal of Global Information Management (JGIM), vol. 31, no. 3, pp. 1-16, 2023, doi: 10.4018/JGIM.324465.

[12] S. Hanumanthaiah, "SOX Considerations for Cloud Data Architecture: A Comprehensive Literature Review," IJSAT-International Journal on Science and Technology, vol. 16, no. 2, pp. 1-11, 2025, doi: 10.71097/IJSAT.v16.i2.6482.

[13] J. H. F. Moreno, G. D. R. Quiñónez, F. G. E. Moreno, and L. A. C. Bone, "Robotic Process Automation (RPA) as a technological tool for automating the execution of audits," Sapienza: International Journal of Interdisciplinary Studies, vol. 4, no. 4, Art. no. e23059-e23059, 2023, doi: 10.51798/sijis.v4i4.658.

[14] Z. Gu, F. Corcoglioniti, D. Lanti, A. Mosca, G. Xiao, J. Xiong, et al., "A systematic overview of data federation systems," Semantic Web, vol. 15, no. 1, pp. 107-165, 2024, doi: 10.3233/SW-223201.

[15] S. Mishra and S. Konidala, "A polyglot data integration framework for seamless integration of heterogeneous data sources and formats," International Journal of Emerging Trends in Computer Science and Information Technology, vol. 5, no. 4, pp. 70-81, 2024, doi: 10.63282/3050-9246.IJETCSIT-V5I4P108.

[16] M. H. Safarzadeh and M. Derakhshan, "Risk Disclosure Quality Assessment Using Hidden Markov Chain Analysis of Firms' Risk State," Computational Economics, pp. 1-44, 2025, doi: 10.1007/s10614-025-10948-7.

[17] J. Lillestøl, "Sampling risk evaluations in tax audits: Some modelling issues," Law, Probability and Risk, vol. 21, no. 1, pp. 1-20, 2022, doi: 10.1093/lpr/mgac010.

[18] L. Gao, "Enterprise internal audit data encryption based on blockchain technology," PLoS one, vol. 20, no. 1, Art. no. e0315759, 2025, doi: 10.1371/journal.pone.0315759.

[19] Y. Qu, H. Ji, and D. Cofell, "Financial Audit Method Innovation of Qinghai Energy Enterprises Based on Panel Data Regression Model," Wireless Communications and Mobile Computing, vol. 2023, no. 1, Art. no. 6093443, 2023, doi: 10.1155/2023/6093443.

[20] Z. Yang, "Privacy-Aware Financial Risk Control: A Federated Learning Approach with Differential Privacy Optimization," Journal of Computer Technology and Software, pp. 4(4), 2025, doi: 10.5281/zenodo.15340786.

[21] I. Namatevs, K. Sudars, A. Nikulins, and K. Ozols, "Privacy auditing in differential private machine learning: The current trends," Applied Sciences, vol. 15, no. 2, pp. 647, 2025, doi: 10.3390/app15020647.

[22] Y. Shin, H. Kim, J. Jeong, and D. Shin, "Federated learning for surveillance systems: A literature review and AHP expert-based evaluation," Electronics, vol. 14, no. 17, pp. 3500, 2025, doi: 10.3390/electronics14173500.

[23] S. A. Mahmud, N. Islam, Z. Islam, Z. Rahman, and S. T. Mehedi, "Privacy-preserving federated learning-based intrusion detection technique for cyber-physical systems," Mathematics, vol. 12, no. 20, pp. 3194, 2024, doi: 10.3390/math12203194.

[24] X. Gu, F. Sabrina, Z. Fan, and S. Sohail, "A review of privacy enhancement methods for federated learning in healthcare systems," International Journal of Environmental Research and Public Health, vol. 20, no. 15, pp. 6539, 2023, doi: 10.3390/ijerph20156539.

[25] V. Feldman and T. Zrnic, "Individual privacy accounting via a Renyi filter," Advances in Neural Information Processing Systems, vol. 34, pp. 28080–28091, 2021, doi: 10.48550/arXiv.2008.11193.

[26] M. Ghazali, D. B. Nugroho, and Y. Latief, "Analyzing the Effect of the Construction Safety Audit Model Using the Business Process Model and Notation (BPMN) Method on Improving Communication and Collaboration Between Stakeholders," CSID Journal of Infrastructure Development, vol. 7, no. 2, pp. 8, 2024, doi: 10.7454/jid.v7.i2.1144.

[27] W. Li, "Audit automation process and realization path analysis based on financial technology," European Journal of Business, Economics & Management, vol. 1, no. 2, pp. 69-75, 2025, [Online]. Available: https://ideas.repec.org/a/dba/ejbema/v1y2025i2p69-75.html.

[28] S. Anvarkhatibi, H. R. Baradaran, A. Mottaghi, and H. Taghizadeh, "Recognition and ranking the effective factors on audit quality via the TOPSIS technique," Advances in Mathematical Finance & Applications, vol. 9, no. 1, pp. 159-180, 2024, doi: 10.22034/AMFA.2022.1966676.1790.

[29] K. F. Liew, W. S. Lam, and W. H. Lam, "Financial network analysis on the performance of companies using integrated entropy–DEMATEL–TOPSIS model," Entropy, vol. 24, no. 8, pp. 1056, 2022, doi: 10.3390/e24081056.