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

A Cross-Domain Data Flow Security Governance and Privacy Protection Mechanism for Distributed Digital Art Platforms

Meixuan Lv[1, *]

[1]School of Art & Design, Wuhan Institute of Technology, Wuhan, 430205, China

Abstract

INTRODUCTION: Distributed digital art platforms face ownership ambiguity, privacy leakage, unauthorized access, and audit difficulties in cross-domain data flows.
OBJECTIVES: To enhance trust, privacy, controllability, and auditability through a dedicated data security governance mechanism..
METHODS: A closed-loop governance mechanism jointly models art data, AIGC models, copyrights, and transactions, integrating five coordinated modules: damage-tolerant ownership verification, privacy-preserving cross-chain auditing, dynamic access control, secure delivery with rollback, and efficient traceability.
RESULTS:Ownership confirmation reaches 97.9%–98.7% (mild perturbations) and 85.9% (combined); authorization accuracy: 95.6%; illegal access interception: 94.9%; normal transaction audit pass rate: 99.3%; anomaly detection: up to 100%; on-chain storage compression: 99.99% (50 MB); traceability success: 94.8%. Access latency remains below 20 ms with 180+ transactions/s under concurrency.
CONCLUSION: The mechanism effectively secures cross-domain data flows with high trustworthiness and low overhead, suitable for distributed digital art ecosystems.

Keywords: cross-domain data flow; distributed system security; privacy protection; dynamic access control; cross-chain audit; digital art platform

Received on 01 June 2026, accepted on 14 June 2026, published on 23 June 2026

doi: 10.4108/eetsis.13268

1. Introduction

1.1. Research Background and Significance

With the integrated application of cloud computing, blockchain, artificial intelligence, and AIGC technologies, digital art design platforms are evolving from single content-production tools into distributed data circulation environments in which multiple subjects, platforms, and blockchain systems collaborate. In this environment, data such as original art assets, prompts, style parameters, user behaviors, AIGC model weights, copyright credentials, and transaction records continuously flow among creators, platforms, model service providers, copyright agencies, traders, and regulators, forming typical cross-domain data flows in emerging distributed networks and systems.

Compared with traditional content management systems, data flows in distributed digital art platforms are characterized by multimodality, high value, strong copyright relevance, high privacy sensitivity, and easy replicability. Without effective governance during cross-domain circulation, privacy inference, permission abuse, model asset leakage, data fraud, transaction repudiation, and difficulties in copyright accountability may occur.

*Corresponding author. Email: 3436193927@qq.com

These risks are highly consistent with data leakage, ambiguous permission boundaries, unverifiable cross-chain sharing, and distributed AI model security issues in emerging distributed systems.

Therefore, this paper takes distributed digital art platforms as a representative application scenario and studies cross-domain data flow security governance and privacy protection mechanisms. The objective is to realize trusted subjects, controllable permissions, protected content, verifiable transactions, and accountable responsibility while enabling data circulation and utilization, thereby providing a reusable technical framework for data security and privacy protection in distributed networks and systems.

1.2. Related Work

Existing research on data security and privacy protection in distributed networks and systems mainly focuses on data asset circulation, digital content generation, cross-domain data flow governance, privacy-preserving computation, cross-chain auditing, access control, and distributed AI model security.

Regarding data markets and the foundations of digital art generation, studies on data markets, data pricing, and personal data trading provide theoretical support for data asset circulation, platform-based transactions, and privacy governance [1–4]. Neural style transfer, vision-language pre-training, diffusion models, and latent diffusion models have promoted automated generation and cross-platform dissemination of digital art content [5–8]. These studies show that data in distributed digital art platforms has both asset circulation attributes and generative AI model service attributes.

In terms of cross-domain data flow security governance, blockchain, distributed storage, and fine-grained access control have been widely used to enhance transparency, tamper resistance, and accountability in data sharing [9–11]. However, existing schemes are mostly designed for general data assets or single-chain evidence storage scenarios and lack unified modeling for heterogeneous objects such as prompts, generated content, model assets, and copyright credentials in digital art data.

In terms of identity authentication and dynamic access control, DID/VC, XACML, ABAC, and UCON can support cross-domain subject identification, attribute verification, and continuous usage control [12–16]. Nevertheless, static policies are difficult to adapt to dynamic changes in subject reputation, access purposes, operation risks, and contextual states, and they also cannot effectively support cross-domain authorization inheritance and policy collaboration.

In terms of privacy-preserving computation and distributed AI security, differential privacy, federated learning, secure aggregation, and secure multi-party computation can support "data usable but invisible" [17–22]. However, in AIGC model services and high-frequency interactive design scenarios, existing solutions still face high computational overhead, large communication costs, insufficient adaptation to model inference, and degraded user experience.

In terms of cross-chain auditing and transaction accountability, commitment mechanisms, secret sharing, zero-knowledge proofs, and smart contracts can verify transaction authenticity and prevent payment repudiation without disclosing data plaintext [23–25]. However, for off-chain large-scale art data, model files, and generated content, low-overhead evidence storage, integrity delivery verification, and arbitrable rollback mechanisms remain insufficient.

In terms of distributed AI model asset protection, DNN model watermarking and deep neural network watermarking techniques can be used for model copyright declaration, integrity verification, and infringement tracing [26–27]. However, in digital art platforms, model watermarking, data watermarking, cross-chain auditing, and dynamic access control remain insufficiently coordinated, making it difficult to support cross-domain asset governance for AIGC model services.

In summary, existing research has not yet formed a closed-loop security mechanism for cross-domain data flows in distributed digital art platforms. On the one hand, integrated modeling of multimodal art data, AIGC model services, and copyright credentials is lacking; on the other hand, identity authentication, access control, privacy-preserving transactions, cross-chain auditing, and infringement traceability are still fragmented. Accordingly, this paper constructs an integrated mechanism that combines privacy protection, dynamic authorization, cross-chain auditing, and regulatory arbitration.

Recent studies on AIGC watermarking, generative content provenance, cross-chain interoperability, privacy-preserving identity authentication, and model watermarking indicate that single techniques can support isolated functions such as provenance labeling, source verification, or privacy authentication [28–35]. However, these studies usually do not jointly bind data ownership attributes, transaction commitments, access-control states, cross-chain audit evidence, and watermark-based traceability into one enforceable lifecycle governance loop. The specific novelty of this paper therefore lies not in replacing mature cryptographic primitives, but in designing an integrated, reproducible, and arbitrable mechanism that coordinates these primitives for cross-domain digital art data circulation.

More specifically, recent work on zero-knowledge data auditing for blockchain-based decentralized storage has demonstrated the feasibility of combining retrievability with verifiable auditability [36]. Lightweight and semi-asynchronous federated learning studies have further shown that distributed AI services can reduce communication overhead and improve collaboration in dynamic networks [37,38]. Blockchain-enabled network security defense in SDN-based IIoT also supports the use of tamper-resistant ledgers for distributed security governance [39]. In addition, game-theoretic MEC

offloading research provides useful theoretical support for adaptive decision-making under dynamic resource and risk constraints [40]. These studies provide direct technical references for the transaction audit, privacy-preserving computation, distributed security governance, and dynamic authorization modules of this paper.

1.3. Research Content and Main Innovations

To address the practical problems of diverse data types, complex subject relationships, dynamically changing authorization boundaries, and prominent privacy leakage risks in distributed digital art platforms, this paper studies security governance and privacy protection mechanisms for cross-domain data flows. First, a cross-domain data flow model for distributed digital art platforms is constructed, and objects such as original assets, user interaction data, AIGC-generated content, model parameters, copyright credentials, and transaction records are classified and graded. Second, a security governance framework integrating decentralized identity, verifiable credentials, and risk-aware dynamic access control is designed. Third, privacy-preserving computation, secret sharing, zero-knowledge proofs, and cross-chain auditing are introduced to achieve "data usable but invisible," verifiable transactions, and traceable accountability. Finally, digital watermarking and model watermarking are combined to support copyright declaration, integrity verification, and infringement tracing for digital artworks and distributed AI model assets. The overall framework is illustrated in Figure 1.

The distinguishes existing technologies from the original contribution of this paper are shown as table1 :

Table 1. Clarification of technical novelty compared with existing schemes

Component	Existing focus	Original contribution in this paper
Blockchain copyright/evidence storage	Single-chain hash evidence and transaction records	Cross-consortium-chain audit with commitment, receipt proof, rollback state, and regulatory arbitration evidence
Watermark/provenance methods	Content source labeling or model watermark detection	Multi-evidence ownership verification combining content, semantic, provenance, watermark, metadata, and on-chain certificate evidence
Static RBAC/ABAC/UCON	Rule-based or attribute-based authorization	Game-theoretic trust prediction and threshold updating for operation-sensitive dynamic authorization
Privacy-preserving transactions	Isolated ZKP or secret-sharing protocols	Scenario-adaptive transaction audit linking off-chain encrypted delivery, on-chain commitment, receipt proof, and payment settlement
Digital asset governance frameworks	Conceptual lifecycle management	Executable lifecycle control points, algorithms, prototype parameters, scalability tests, and limitations analysis

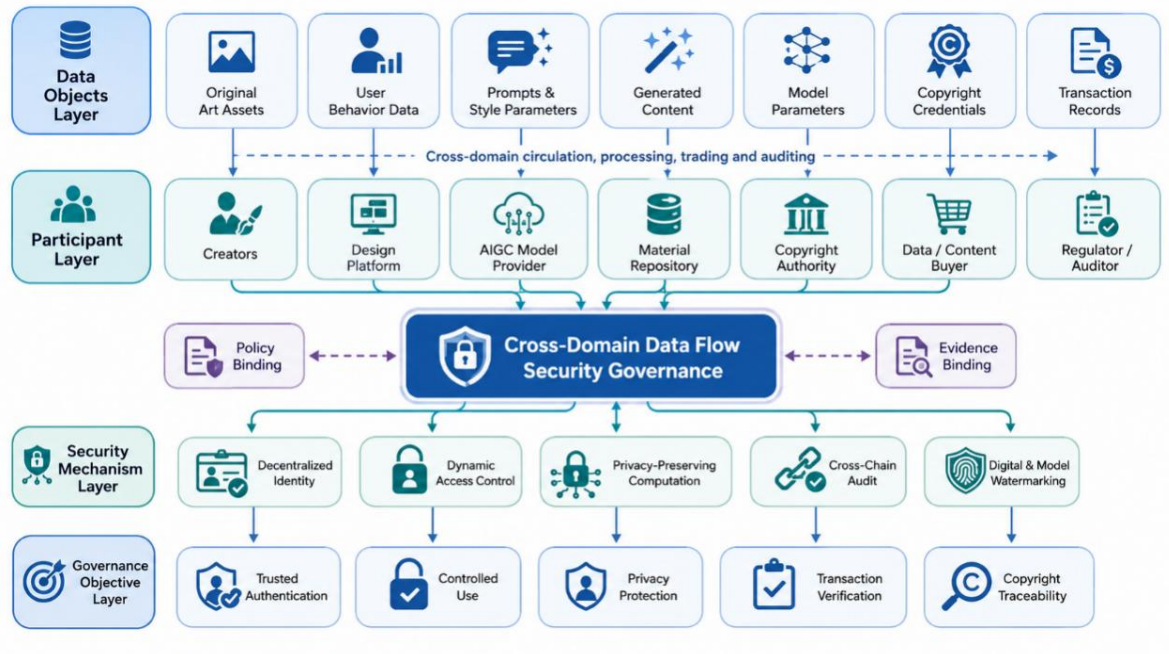

Figure 1. Overall research framework for cross-domain data flow security governance in distributed digital art

(1) A cross-domain data flow model for distributed digital art platforms is constructed to uniformly describe multimodal art data, AIGC model services, copyright credentials, and transaction records;
(2) A security governance framework that collaboratively binds identity, data, policies, and evidence is developed to improve the trustworthiness, controllability, and auditability of cross-platform data circulation;
(3) A risk-aware dynamic access control mechanism is designed by introducing subject trustworthiness, data sensitivity levels, operation risks, and contextual states to realize adaptive authorization;
(4) Privacy-preserving computation, zero-knowledge proofs, secret sharing, and cross-chain auditing are integrated to support transaction verification and regulatory arbitration while protecting data content privacy;
(5) Low-overhead evidence storage, multi-evidence violation traceability, and digital/model watermarking are introduced to strengthen copyright accountability for distributed AI model assets and digital art content.

2. System Model and Technical Foundations

2.1 Cross-Domain Data Flow Model for Distributed Digital Art Platforms

Data flows in distributed digital art platforms exhibit multimodality, cross-domain collaboration, strong privacy sensitivity, and strong rights-related attributes. Data objects include original creative assets, process-derived data, AIGC-generated content, model parameters, copyright credentials, and transaction records; data subjects include creators, platforms, model service providers, material providers, copyright agencies, trading parties, and regulators. The circulation process can be abstracted as a chain of "data generation—platform processing—model invocation—authorization transaction—content publication—audit and accountability."
As shown in Figure 2, digital art data flows across content domains, model service domains, transaction domains, and regulatory domains in a distributed platform. This process must support the release of data value while continuously maintaining trusted identity, policy consistency, privacy protection, and verifiable evidence. DID authentication, dynamic access control, privacy-preserving computation, cross-chain evidence storage, and watermark-based tracing jointly constitute the key support for cross-domain circulation.

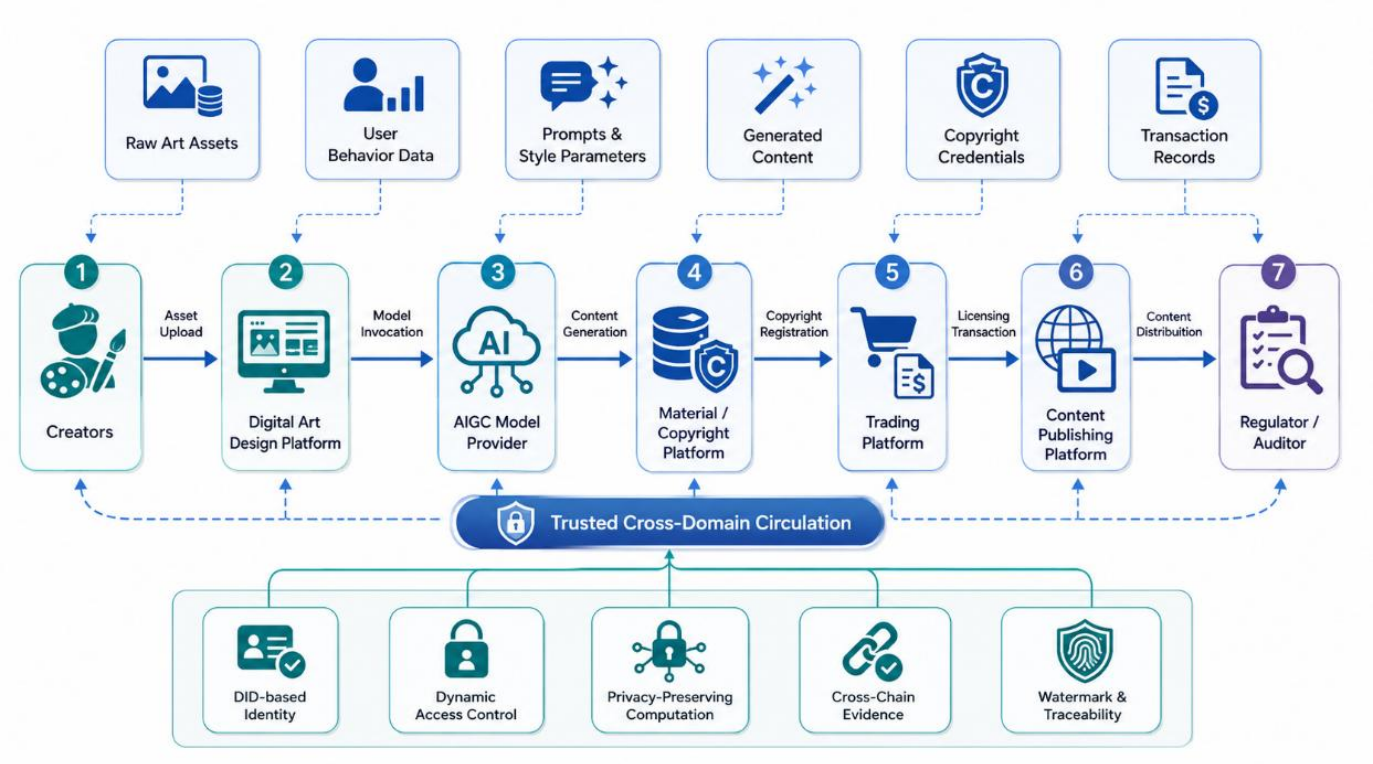

Figure 2. Cross-domain circulation model of digital art data in a distributed platform

Let the set of data objects in the digital art platform be D, the set of subjects be S, and the set of operations be O. The cross-domain data flow can then be represented as a directed relation composed of subjects, data, operations, time, scenarios, and permission constraints, as shown in Eq. (1):

$$F = (S, D, O, C, P, E) \tag{1}$$

In Eq. (1), S denotes the subject set, D denotes the data object set, O denotes the operation set, C denotes contextual scenarios and domain constraints, P denotes access and usage policies, and E denotes audit and evidence information. This model can uniformly describe the circulation, authorization, and governance logic of digital art data across platforms, institutions, and domains, thereby providing a basic model for security mechanism design.

It can therefore be seen that the core contradiction in distributed digital art platforms lies in the fact that data value depends on open circulation, whereas security governance depends on cross-domain boundary control. The subsequent mechanisms in this paper are developed around "data usability, privacy protection, transaction verifiability, and accountability."

2.2. Security Objectives and Theoretical Foundations of Cross-Domain Data Flows

Data element circulation emphasizes asset-based flow in stages such as ownership confirmation, registration, transaction, delivery, use, and destruction. When mapped to distributed digital art platforms, the security objectives can be summarized as controllability, verifiability, and traceability: controllability requires that data be used according to authorization policies; verifiability requires that sources, transactions, and invocation processes be verifiable; and traceability requires that leakage, infringement, and violations be locatable.

Accordingly, this paper adopts "prior ownership confirmation and classification, in-process dynamic governance, and post-event audit and arbitration" as the overall security logic, and forms a full-lifecycle data governance loop through policy flows, evidence flows, and risk feedback flows.

2.3 Cross-Chain Transaction Auditing and Privacy Protection Technologies

Cross-chain technology is used to realize trusted message transmission, state verification, and transaction collaboration among different consortium chains or blockchain systems. For distributed digital art platforms, this paper adopts an architecture combining relay chains and consortium chains, and integrates Pedersen commitments, Shamir secret sharing, zero-knowledge proofs, and smart contracts to verify transaction authenticity, delivery integrity, and payment fairness without disclosing raw data.

2.4. Decentralized Identity and Dynamic Access Control Technologies

Decentralized identity (DID) and verifiable credentials (VC) can provide verifiable identity and attribute proofs for cross-domain subjects. Dynamic access control adjusts authorization policies in real time according to subject reputation, data sensitivity levels, operation risks, and contextual states. Their combination can form a closed loop of "identity authentication—attribute verification—risk assessment—dynamic authorization—behavior feedback," which is suitable for continuous governance of cross-domain data flows in distributed environments.

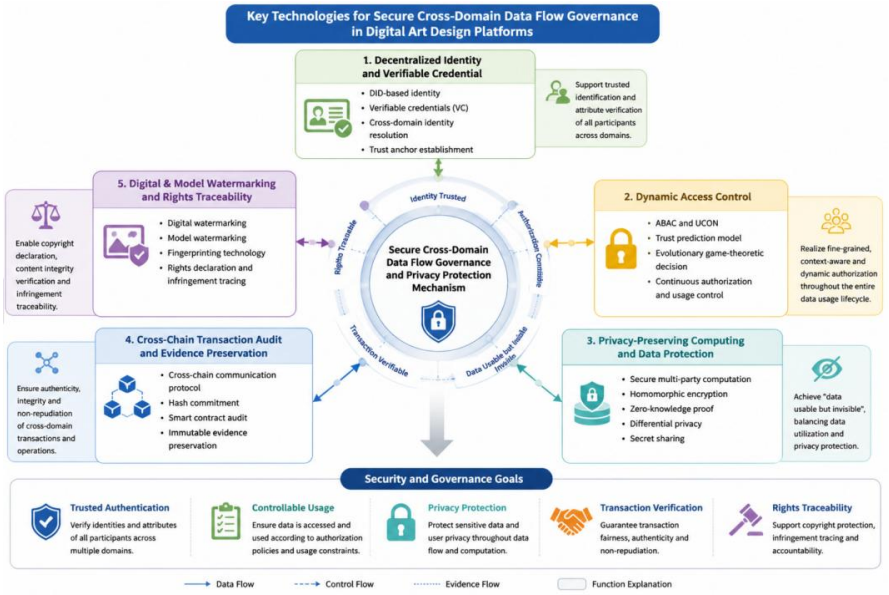

Figure 3. Relationship among key technologies for secure cross-domain data flow governance in distributed systems

As shown in Figure 3, cross-domain data flow security governance is supported by multiple key technologies. DID/VC solves trusted subject identification; dynamic access control realizes fine-grained authorization; privacy-preserving computation supports data usability without visibility; cross-chain transaction auditing ensures transaction verifiability; and digital watermarking and model watermarking enable copyright accountability. These technologies jointly correspond to the topics of data security, privacy-preserving computation, cross-domain flow control, and distributed AI model asset protection in distributed networks and systems.

2.5. Limitations of Existing Schemes

Existing schemes still have several limitations for distributed digital art platforms: cross-domain data flow models are insufficiently integrated with the semantics of art data; transaction auditing for off-chain large files and model assets is difficult to verify at low cost; identity, data, authorization, and evidence are not strongly bound; static access control cannot adapt to dynamic risks; and privacy protection and copyright accountability mechanisms lack coordination. Therefore, an integrated security governance mechanism for distributed system scenarios is required.

3. Cross-Domain Data Flow Security Governance Framework for Distributed Digital Art Platforms

To address diverse data assets, complex subject relationships, frequent cross-domain circulation, and unclear security responsibilities in distributed digital art platforms, this paper constructs a cross-domain data flow security governance framework. Taking full-lifecycle data governance as the main line, the framework uses data asset identification, decentralized identity, risk-aware access control, privacy-preserving computation, cross-chain transaction auditing, and watermark-based tracing as core mechanisms to achieve trusted authentication, controllable use, privacy protection, transaction verification, and copyright traceability.

Unlike traditional data sharing frameworks, the proposed framework targets a data circulation environment involving cross-platform collaboration, cross-chain interaction, and distributed AI model services. Therefore, it does not treat a digital art platform as a single application system, but abstracts it as a representative data security scenario in emerging distributed networks and systems.

3.1. Design Objectives, Threat Assumptions, and Security Principles

For distributed digital art platforms, the proposed framework takes data value release and controllable security risks as its core and sets the following objectives:

(1) Trusted subject authentication: establish cross-domain identity identifiers and attribute verification mechanisms for creators, platforms, model service providers, copyright agencies, and regulatory nodes;

(2) Identifiable data assets: uniformly identify and bind policies to materials, prompts, generated content, models, copyright credentials, and transaction records;

(3) Controllable data circulation: propagate classification, grading, access policies, and usage constraints with data across domains to prevent unauthorized access and illegal dissemination;

(4) Privacy-preserving data usability: use privacy-preserving computation to support secure utilization of sensitive data in model inference, feature computation, and collaborative analysis;

(5) Verifiable transaction process: rely on cross-chain evidence storage, commitment mechanisms, zero-knowledge proofs, and smart contracts to prevent transaction fraud and repudiation and to support auditing;

(6) Traceable copyright responsibility: integrate digital watermarking, model watermarking, and evidence-chain tracing to realize infringement identification and responsibility tracing.

In an open cross-domain environment, the system mainly faces threats such as identity forgery and unauthorized access, data fraud and payment repudiation, privacy inference by semi-honest service providers, model asset abuse, evidence non-mutual recognition, and cross-chain state inconsistency. This paper assumes that underlying cryptographic components are secure, consortium-chain ledgers are tamper-resistant under consensus constraints, smart contracts execute according to predefined logic, and no large-scale collusion exceeding the threshold occurs.

The threat model further assumes five adversarial behaviors: A1 identity forgery or credential replay by an external attacker; A2 semi-honest but curious model service providers attempting to infer prompts, input features, or commercial design intentions; A3 malicious buyers attempting payment repudiation, unauthorized re-dissemination, or out-of-scope model training; A4 dishonest data sharers submitting forged or tampered data; and A5 partial cross-chain desynchronization caused by delayed state synchronization. Collusion is bounded by the threshold parameter t in the secret-sharing module, and at least one consortium-chain audit node in each domain is assumed to remain honest. These assumptions are used consistently in the security analysis in Section 5.4.4.

3.2. Overall Architecture

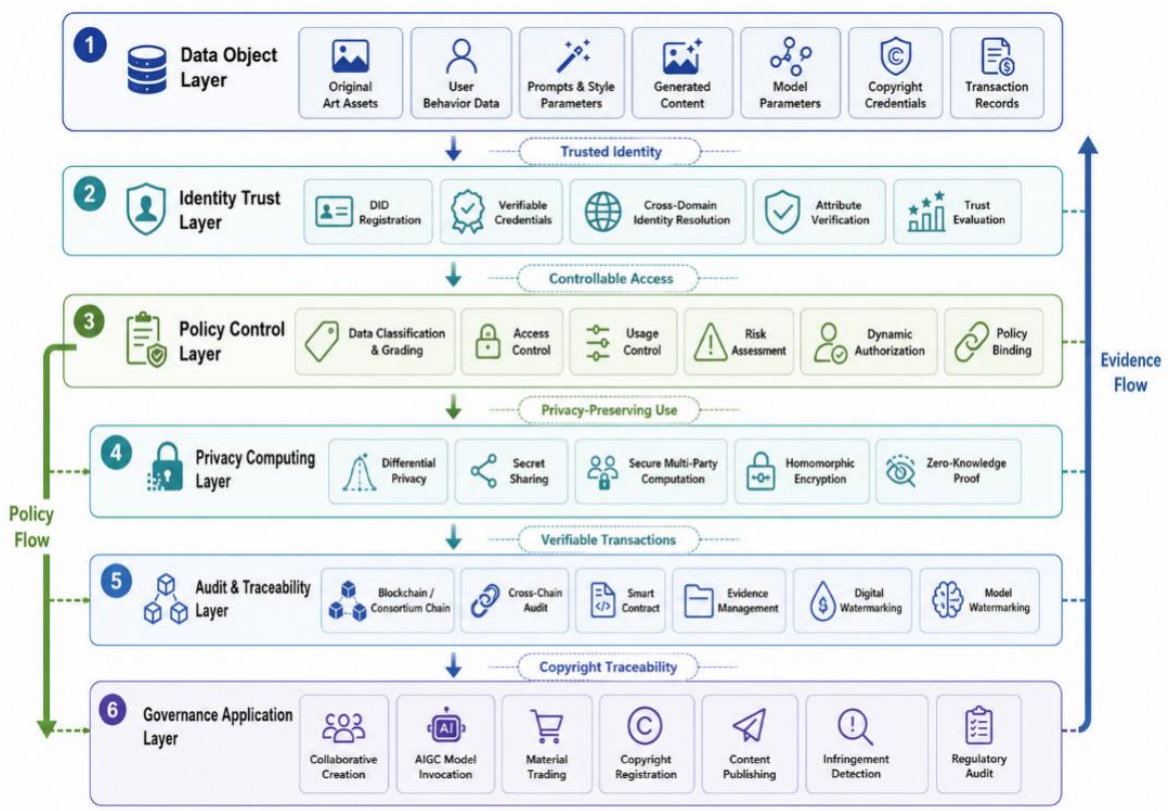

Figure 4. Overall architecture of the cross-domain data flow security governance framework for distributed digital art platforms

The proposed cross-domain data flow security governance framework adopts a layered structure consisting of the data object layer, identity trust layer, policy control layer, privacy computing layer, audit and traceability layer, and governance application layer. To highlight its adaptability to distributed systems, Figure 4 shows the collaboration among these layers in cross-domain data flows, cross-chain state synchronization, and privacy-preserving computation.

As shown in Figure 4, the data object layer is responsible for identification, classification, and grading of multimodal art data and model assets; the identity trust layer performs cross-domain subject authentication based on DID/VC; the policy control layer realizes dynamic authorization and continuous usage constraints; the privacy computing layer supports data usability without visibility; the audit and traceability layer forms an evidence chain through on-chain evidence storage, cross-chain auditing, and watermark-based tracing; and the governance application layer provides security services for collaborative creation, AIGC model invocation, material trading, copyright registration, content publication, and regulatory auditing.

Through the linkage of policy flows and evidence flows, this architecture supports verifiable subjects, controllable data, protected privacy, verifiable transactions, and accountable responsibility in cross-domain circulation of distributed digital art data.

3.3. Full-Lifecycle Security Process of Cross-Domain Data

Cross-domain circulation of digital art data runs through the full lifecycle of generation, registration, authorization, computation, transaction, audit, and destruction. This paper divides it into seven stages and sets identity, policy, privacy, and evidence control points at each stage.

Cross-domain circulation of digital art data is not a single data transmission behavior, but a full-lifecycle process covering data generation, registration, classification, authorization, computation, transaction, publication, audit, and destruction. This paper abstracts it into seven consecutive stages and sets corresponding security control points at each stage. To further illustrate the state transition relationship and security control logic of digital art data during cross-domain circulation, Figure 5 presents the full-lifecycle security process of cross-domain digital art data flow.

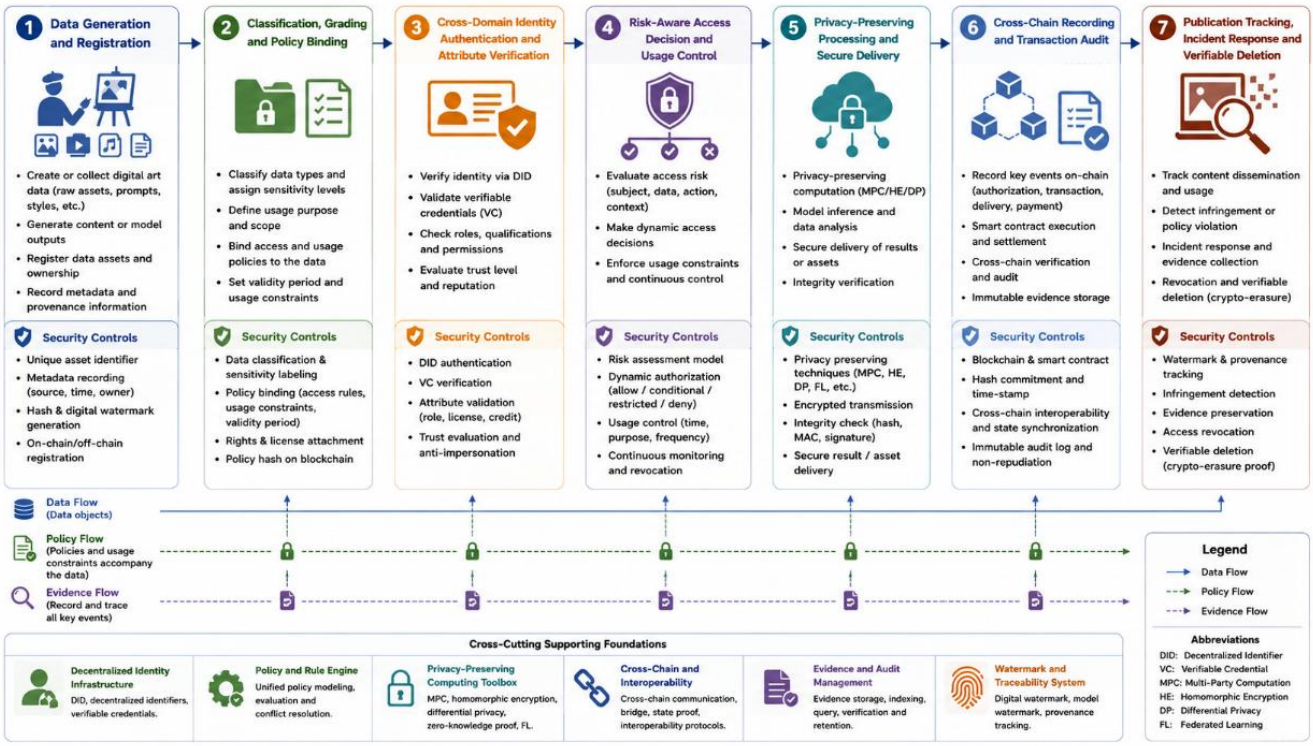

Figure 5. Full lifecycle security process of cross-domain digital art data flow

As shown in Figure 5, the data generation and registration stage emphasizes asset identification and ownership records; the classification and grading stage emphasizes sensitivity levels and policy binding; the identity authentication stage emphasizes DID/VC verification; the access decision stage emphasizes risk-aware dynamic authorization; the privacy processing stage emphasizes secure computation and encrypted delivery; the transaction auditing stage emphasizes cross-chain evidence storage and non-repudiation; and the publication tracking stage emphasizes watermark-based traceability, violation response, and verifiable deletion.

To formally describe this lifecycle process, let the digital art data object be di, and let its lifecycle state set be, as shown in Eq. (2):

$$L_i = l_1, l_2, \ldots, l_7 \tag{2}$$

In Eq. (2), l1 to l7 denote generation and registration, classification and grading, identity verification, access decision, privacy processing, transaction auditing, and publication tracking, respectively.

For any state transition, the system should satisfy, as shown in Eq. (3):

$$T(l_k, l_{k+1}) \Rightarrow Auth(s, d_i, o, c) \wedge Policy(d_i) \wedge Evidence(e) \tag{3}$$

In Eq. (3), Auth(s, di, o, c) denotes the authorization result of subject s performing operation o on data di under context c; Policy(di) denotes the access and usage policy bound to the data object; and Evidence(e) denotes the audit evidence generated by the state transition. This constraint indicates that each cross-domain circulation of digital art data must simultaneously satisfy valid authorization, policy consistency, and evidence verifiability.

3.4. Core Mechanisms and Their Collaboration

To support the above framework, this paper designs five core mechanisms and realizes their collaboration through policy flows, evidence flows, and risk feedback flows.

(1) Data asset identification and policy binding mechanism. This mechanism generates a unique asset identifier for each digital art data object and writes data type, sensitivity level, ownership information, authorization scope, usage purpose, and circulation rules into its metadata description. For tradable or authorizable data, the system further binds contract terms, copyright credentials, and watermark identifiers. This mechanism addresses the questions of "what the data is, who owns it, and how it may be used," and serves as the foundation for subsequent identity verification, access control, and audit accountability.

(2) DID/VC-based cross-domain identity trust mechanism. This mechanism establishes decentralized identities for subjects on different platforms and expresses subject roles, platform qualifications, credit levels, copyright registration qualifications, and transaction permissions through verifiable credentials. When a subject initiates an access request, the system resolves

verification keys and service endpoints through DID, verifies subject attributes through VC, and inputs the verification result into the access control module. This mechanism addresses "who the subject is, whether it is trustworthy, and whether it has access eligibility."

(3) Risk-aware dynamic access control mechanism. This mechanism makes access decisions by integrating subject risk, data risk, operation risk, and contextual risk. Suppose subject s performs operation o on data d. The comprehensive risk score is defined as:

$$R(s,d,o,c) = \alpha R_s + \beta R_d + \gamma R_o + \delta R_c \quad (4)$$

where Rs denotes subject risk, Rd denotes data sensitivity risk, Ro denotes operation risk, and Rc denotes contextual risk, with the parameters serving as weights. When the risk score is below the first threshold, the system allows access; when it falls between thresholds, the system performs conditional or degraded access; and when it exceeds the upper threshold, the system denies access and triggers auditing. This mechanism addresses whether access should be allowed, to what extent it should be allowed, and whether additional constraints are required.

(4) Privacy-preserving computation and secure delivery mechanism. This mechanism selects appropriate privacy-preserving technologies according to business scenarios and data sensitivity levels. Secure multi-party computation or secret sharing can be used in cross-domain feature analysis; differential privacy can be used in user behavior statistics; encrypted inference or secure protocols can be used in model service invocation; and commitment mechanisms and zero-knowledge proofs can be used in transaction verification. This mechanism addresses how data can be utilized without exposing the raw data.

(5) Cross-chain auditing and copyright tracing mechanism. This mechanism writes transaction records, authorization states, data digests, model invocation logs, and copyright credentials into on-chain or cross-chain audit systems, and combines digital watermarking and model watermarking to realize infringement tracing. Cross-chain auditing addresses non-mutual recognition of evidence between different platforms, while watermark-based tracing addresses copying, tampering, and illegal dissemination after data leaves its original domain. This mechanism addresses whether a transaction is authentic, whether authorization is valid, and whether infringement can be held accountable.

The above mechanisms form a closed-loop collaboration through three types of relationships: policy flows carry access policies and usage constraints across domains with data; evidence flows record identity verification, authorization decisions, privacy-preserving computation, transaction fulfillment, and copyright tracing results; and risk feedback flows dynamically update subject reputation and authorization thresholds according to violation events.

Therefore, the proposed framework unifies subject trust, policy control, privacy protection, transaction auditing, and copyright accountability into one governance loop, providing scenario-oriented support for cross-domain data flow security in distributed networks and systems.

4. Cross-Domain Ownership Verification, Privacy-Preserving Transactions, and Dynamic Access Control Mechanisms

Based on the overall framework in Section 3, this section further presents the key mechanism design for distributed digital art platforms. Digital art data and AIGC model assets are easy to copy, tamper with, disseminate, and involve complex ownership boundaries. Relying solely on static evidence storage or traditional access control cannot satisfy ownership verification, transaction, authorization, and accountability requirements in cross-domain circulation.

Let the digital art data object be d, the data owner be Uo, the data buyer or user be Ub, the ownership confirmation authority be CA, the transaction audit contract set be SC, and the cross-domain chain system include the ownership-domain consortium chain Bo, the usage-domain consortium chain Bb, and the relay chain Br. The metadata of data object d is represented as, as shown in Eq. (5):

$$M_d = (ID_d, \mathrm{Type}_d, \mathrm{Owner}_d, \mathrm{Level}_d, \mathrm{Policy}_d, \mathrm{Hash}_d, \mathrm{Cert}_d, \mathrm{Trace}_d) \quad (5)$$

In Eq. (5), IDd denotes the data asset identifier, Typed denotes the data type, Ownerd denotes the ownership subject, Leveld denotes the sensitivity level, Policyd denotes the access and usage policy, Hashd denotes the data digest, Certd denotes the ownership certificate, and Traced denotes the circulation tracing information. All mechanisms in this section are developed around this metadata structure.

4.1. Arbitrable Damage-Tolerant Data Ownership Verification Method

The core issue of data ownership verification is that when a work or model output is compressed, cropped, format-converted, locally tampered with, style-transferred, or regenerated, the system should still be able to identify its ownership source and provide arbitrable evidence in dispute scenarios. Unlike single-hash ownership verification, this paper integrates content features, semantic features, provenance features, watermark identifiers, and on-chain certificates to improve the robustness of data ownership verification in distributed environments.

For a digital art data object d, the system first extracts its ownership attributes from multiple dimensions. Let the ownership attribute set be, as shown in Eq. (6):

$$A_d = \{A_c, A_s, A_p, A_w, A_m\} \quad (6)$$

In Eq. (6), Ac denotes content visual attributes, including color distribution, texture features, edge structures, and local key points; As denotes semantic attributes, including

subject tags, object categories, style descriptions, and text-prompt semantic representations; Ap denotes provenance attributes, including creator identity, creation time, device fingerprint, platform identifier, and upload record; Aw denotes watermark attributes, including explicit watermarks, implicit watermarks, and model-generated fingerprints; and Am denotes metadata attributes, including file format, layer structure, copyright statement, and authorization policy.

Attributes that can uniquely identify the ownership subject, such as creator DID, digital signature, platform upload record, and copyright registration number, may be regarded as special attributes. Attributes that cannot independently confirm ownership but provide auxiliary evidence, such as image style, content semantics, prompts, and editing traces, are regarded as general attributes and must participate in ownership verification in combination. Thus, the system obtains the ownership feature vector, as shown in Eq. (7):

$$V_d = \text{Extract}\left(A_c, A_s, A_p, A_w, A_m\right) \tag{7}$$

To avoid directly exposing artwork content and creative privacy, the system does not submit raw data directly to the ownership confirmation authority. Instead, it generates an arbitration code based on the ownership feature vector. The authority CA generates a random collaborative damage-resistant factor rd for the data object and returns usage rules Ruled to the owner. The owner combines the ownership feature vector Vd, random factor rd, and data digest Hashd according to the rules to generate the arbitration code, as shown in Eq. (8):

$$\text{Code}_d = H\left(V_d \| r_d \| \text{Hash}_d \| \text{Owner}_d\right) \tag{8}$$

In Eq. (8), H() denotes a secure hash function and || denotes concatenation. The authority checks whether Coded already exists in the ownership database or on-chain registration system. If duplicate registration is not found, an ownership certificate is issued for the data object, as shown in Eq. (9):

$$\text{Cert}_d = \text{Sign}_{CA}\left(ID_d, \text{Owner}_d, \text{Code}_d, \text{Policy}_d, T_d\right) \tag{9}$$

In Eq. (9), Td denotes the certificate validity period or registration timestamp. After the certificate is issued, the system writes IDd, Coded, the certificate digest, and the timestamp into the on-chain evidence storage system, while the raw art data and complete ownership features are stored in an off-chain secure storage environment.

During cross-domain circulation, digital art data may suffer attribute damage due to compression, cropping, overlaying, format conversion, or local tampering. To support damage-tolerant ownership verification, this paper adopts a method combining similarity measurement and cross-evidence verification. For the data d' to be arbitrated, the system extracts its ownership feature vector and compares it with the ownership feature vector Vd of the originally registered data d, as shown in Eq. (10):

$$\text{Sim}\left(d, d'\right) = \lambda_1 \text{Sim}_c + \lambda_2 \text{Sim}_s + \lambda_3 \text{Sim}_p + \lambda_4 \text{Sim}_w + \lambda_5 \text{Sim}_m \tag{10}$$

In Eq. (10), Simc, Sims, Simp, Simw, and Simm denote the similarities at the content, semantic, provenance, watermark, and metadata levels, respectively; λi denotes the weight parameter and satisfies the corresponding normalization constraint.

If the similarity satisfies the ownership threshold, the system determines that d' is associated with the registered data d; if it is below the threshold, the ownership association is deemed insufficient; if the similarity falls in a critical interval, manual arbitration or multi-evidence cross-verification is initiated. This method avoids the excessive sensitivity of single-hash ownership verification to minor data changes and is suitable for derivative design, local editing, and AIGC regeneration scenarios in digital art works.

When digital art data is authorized, transferred, inherited, or traded, the ownership confirmation system must synchronously update the ownership state. This paper divides ownership changes into three types: adding ownership subjects, which adds co-right holders based on the original owner; replacing ownership subjects, which changes the original owner to a new owner; and copy authorization, in which the original owner retains ownership but grants the user permissions for copying, display, training, generation, or commercial use.

For any ownership change event Et, the system generates an rights-change record, as shown in Eq. (11):

$$\text{Update}_d = \text{Sign}_{CA}\left(ID_d, \text{OldOwner}, \text{NewOwner}, \text{RightType}, \text{Policy}_d', T_t\right) \tag{11}$$

In Eq. (11), RightType denotes the rights type and the policy parameter denotes the updated access and usage policy. The system then invalidates the old certificate or generates a copy authorization certificate and writes the change record into the on-chain evidence storage system. Thus, digital art data can maintain consistent ownership states, verifiable certificates, and traceable responsibility during cross-domain transactions and multiple rounds of authorization.

4.2. Privacy-Preserving Transaction Audit Scheme across Consortium Chains

Digital art data transactions usually occur among different platforms or institutions, each of which may maintain an independent consortium chain or evidence storage system. For example, a material platform records work ownership on consortium chain Bo, the buyer's design platform records transaction requests on consortium chain Bb, and the relay chain Br is responsible for cross-chain state synchronization and contract execution. Because digital art data may contain commercial ideas, unpublished works, user prompts, and model inputs, transaction auditing cannot directly disclose data plaintext. Therefore, this paper designs a privacy-preserving transaction audit scheme across consortium chains to verify transaction authenticity, data integrity, and payment fairness without disclosing data content.

The system includes four types of participants: data sharer Uo, data buyer Ub, ownership-domain consortium chain Bo, usage-domain consortium chain Bb, and relay chain Br. Bo stores data ownership certificates, data digests, and authorization states; Bb records buyer transaction requests and usage policies; and Br deploys the cross-chain audit contract, payment contract, pledge accountability contract, and query contract. The smart contract set is expressed as, as shown in Eq. (12):

$$SC = \{SC_q, SC_a, SC_p, SC_s\} \tag{12}$$

In Eq. (12), SCq denotes the cross-chain query contract for querying data evidence and certificate states; SCa denotes the audit contract for verifying commitments, proofs, and receipts; SCp denotes the payment contract for escrow of buyer funds; and SCs denotes the pledge accountability contract for constraining sharers to submit authentic data and triggering penalties when malicious behavior occurs.

For implementation, SCq exposes queryCertificate(IDd), queryPolicy(IDd), and queryState(IDT); SCa exposes verifyCommitment(CM, Hashd), verifyZKProof(πd), verifyReceipt(πr), and updateAuditState(IDT); SCp exposes escrow(IDT), release(IDT), freeze(IDT), and refund(IDT); and SCs exposes pledgeLock(Uo), penalty(IDT), and blacklistUpdate(s). The contracts use event logs to synchronize transaction states among Bo, Bb, and Br, while off-chain encrypted files are referenced only by digest indexes. This contract-level description was added to make the cross-chain audit procedure independently reproducible.

To ensure that transaction data is not disclosed in plaintext during auditing, the sharer first generates a Pedersen commitment for digital art data d. Let G be a cyclic group of order p, let (g, h) be generators, and let the sharer choose a random number to generate the commitment, as shown in Eq. (13):

$$CM = g^{H(d)} h^r \tag{13}$$

This commitment has hiding and binding properties, and can bind the data digest H(d) without revealing d. Subsequently, the sharer uses Shamir secret sharing to split the data decryption key or data fragment key Kd into n secret shares, as shown in Eq. (14):

$$\text{Share}(K_d) = \{sh_1, sh_2, \ldots, sh_n\} \tag{14}$$

In Eq. (14), any no fewer than t shares can recover Kd, while fewer than t shares reveal no valid information. For large-scale art data, the system does not perform secret sharing on the complete file directly; instead, it secret-shares the data encryption key, and the raw data is delivered off-chain in ciphertext form to reduce computation and communication overhead.

The buyer needs to confirm that the data delivered by the sharer is consistent with the on-chain commitment, but the data content should not be disclosed during auditing. Therefore, the sharer constructs a zero-knowledge proof pid proving knowledge of data d and random number r satisfying the relation in Equation (13). The audit contract verifies, as shown in Eq. (15):

$$\text{Verify}(CM, \pi_d) = 1 \tag{15}$$

If verification succeeds, the data submitted by the sharer is consistent with the commitment; if verification fails, data fraud or commitment inconsistency may exist, and the transaction enters the exception handling process. This process completes data authenticity auditing without revealing digital art data plaintext.

To prevent the buyer from denying receipt after receiving the data, the system introduces a receipt proof mechanism. After receiving encrypted data, secret shares, or decryption materials, the buyer generates receipt proof pir and submits it to the audit contract. The audit contract verifies whether the receipt matches the data commitment, transaction identifier, and buyer identity, as shown in Eq. (16):

$$\text{VerifyReceipt}(\pi_r, CM, ID_T, U_b) = 1 \tag{16}$$

In Eq. (16), ID_T denotes the transaction identifier.

The payment process adopts a "fund escrow—data verification—receipt submission—automatic settlement" procedure. The buyer first locks funds in payment contract SCp; the sharer submits commitments, proofs, and encrypted delivery materials; after data recovery and integrity verification, the buyer submits a receipt proof; if the audit contract passes verification, the payment contract automatically releases funds to the sharer. If verification fails, the funds remain locked or enter the rollback process. Thus, the system can prevent both the sharer from submitting forged data and the buyer from repudiating payment.

The privacy-preserving transaction audit procedure across consortium chains is as follows.

Step 1: The sharer registers data asset d on the ownership-domain consortium chain Bo and submits IDd, Hashd, Certd, and the authorization policy digest.

Step 2: The buyer initiates a transaction request on the usage-domain consortium chain Bb and declares the purchase purpose, usage scope, access period, and payment conditions.

Step 3: Relay chain Br invokes cross-chain query contract SCq to verify the data ownership certificate, authorization status, and transaction qualification.

Step 4: The sharer generates Pedersen commitment CM and submits the commitment and zero-knowledge proof pi_d to audit contract SCa.

Step 5: The sharer encrypts and secret-shares the data or decryption key and delivers it to the buyer through an off-chain secure channel.

Step 6: The buyer verifies data integrity, generates receipt proof pi_r, and submits it to the audit contract.

Step 7: The audit contract verifies the data authenticity proof and receipt proof. If both pass, the payment contract releases funds; otherwise, pledge accountability or transaction rollback is triggered.

Through "on-chain commitment and auditing, off-chain encrypted delivery, cross-chain state synchronization, and automatic contract settlement," the scheme realizes privacy protection and verifiable auditing for cross-platform digital art data transactions.

4.3. Game-Theoretic Trustworthy Dynamic Access Control Mechanism

In digital art platforms, access subjects exhibit uncertain and dynamically changing behaviors. The same subject may present different risks in different scenarios, such as normal browsing, commercial downloading, model training, secondary generation, and batch export. Traditional static access control cannot adjust permissions in real time according to subject trustworthiness and environmental risks. Therefore, this paper proposes a game-theoretic trustworthy dynamic access control mechanism that combines trust prediction, evolutionary games, and incomplete-information games to realize fine-grained dynamic authorization in cross-domain data flows.

Let the trust evaluation attribute set of access subject s be, as shown in Eq. (17):

$$X_s = \{x_1, x_2, \ldots, x_m\} \tag{17}$$

The attributes include identity authentication strength, historical transaction success rate, violation records, access frequency, payment credit, platform level, certificate validity, and behavior anomaly degree. The system assigns a weight wi to each attribute and obtains the subject trust value, as shown in Eq. (18):

$$T_s = \sum_{i=1}^{m} w_i x_i \tag{18}$$

To reflect uncertainty in the trust value, this paper maps Ts to a trust probability, as shown in Eq. (19):

$$p_s = \frac{1}{1 + e^{-\eta(T_s - \mu)}} \tag{19}$$

In Eq. (19), eta is an adjustment coefficient and mu is a trust boundary parameter. A higher ps indicates that the subject is more likely to perform compliant access behavior, whereas a lower ps indicates that the subject is more likely to pose risks such as unauthorized access, illegal downloading, secondary dissemination, or model abuse.

The access control process can be regarded as a game between the access subject and the data controller. The access subject may choose normal access or malicious access, while the data controller may choose to allow, conditionally allow, or deny access. Suppose the comprehensive risk of the subject performing operation o is, as shown in Eq. (20):

$$R(s,d,o,c) = \alpha R_s + \beta R_d + \gamma R_o + \delta R_c \tag{20}$$

In Eq. (20), Rs denotes subject risk, Rd denotes data sensitivity risk, Ro denotes operation risk, Rc denotes contextual risk, and alpha, beta, gamma, and delta are weighting parameters. For digital art data, different operations have different risk intensities. Preview browsing has lower operation risk, whereas downloading original images, batch export, model training, commercial publication, and reauthorization have higher risks. The system constructs the access control utility function according to gains and losses, as shown in Eq. (21):

$$U = P_{\text{gain}} - P_{\text{loss}} - C_{\text{control}} \tag{21}$$

In Eq. (21), Pgain denotes the business gain brought by allowing access, Ploss denotes possible privacy leakage, copyright infringement, or transaction loss, and Ccontrol denotes the security control cost.

The system makes access decisions according to trust probability ps, comprehensive risk R(s, d, o, c), and authorization threshold theta. The authorization function is defined as, as shown in Eq. (22):

$$\text{Decision}(s,d,o,c) = \begin{cases} \text{Allow}, & R(s,d,o,c) \le \theta_1 \wedge p_s \ge p_1 \\ \text{Conditional}, & \theta_1 < R(s,d,o,c) \le \theta_2 \\ \text{Deny}, & R(s,d,o,c) > \theta_2 \vee p_s < p_0 \end{cases} \tag{22}$$

In Eq. (22), Allow denotes access permission, Conditional denotes conditional access, and Deny denotes access denial. Conditional access may include reducing data precision, read-only preview, mandatory watermarking, limiting invocation times, prohibiting downloads, delayed review, or requiring additional deposits.

Access control policies should not remain fixed, but should be dynamically adjusted according to the long-term operational state of the platform. Let x be the proportion of normal accesses and 1-x the proportion of malicious accesses in a certain period. The system updates utility parameters according to access results, violation events, and audit feedback. When malicious access costs are low or violation benefits are high, the system increases penalty factors and authorization thresholds. When subjects remain compliant over a long period, the system reduces access friction and improves service efficiency.

The authorization threshold update rule can be expressed as, as shown in Eq. (23):

$$\theta_{t+1} = \theta_t + \rho\left(V_t - \bar{V}\right) - \sigma\left(G_t - \bar{G}\right) \tag{23}$$

In Eq. (23), Vt denotes the violation rate in the current period, the corresponding barred variable denotes the target violation rate, Gt denotes compliant access benefits, the corresponding barred variable denotes the target benefit, and rho and sigma are adjustment parameters. When the violation rate increases, the system raises the threshold and tightens the access policy; when compliant benefits are stable and the violation rate is low, the system moderately relaxes the policy to improve data circulation efficiency.

After each access ends, the system updates the subject trust value according to access behavior. If the subject uses data according to the authorized purpose, does not perform unauthorized access, and completes transactions on time, the trust value increases; if the subject exhibits abnormal downloading, policy evasion, illegal dissemination, or payment anomalies, the trust value decreases. The trust update function is, as shown in Eq. (24):

$$T_s^{t+1} = \omega T_s^t + (1-\omega) F_s^t \tag{24}$$

In Eq. (24), the feedback term denotes the behavior feedback score in the current period, and omega denotes

the historical trust retention coefficient. Through the trust feedback mechanism, the system forms a closed loop of "trust prediction—access decision—behavior auditing—trust updating," realizing dynamic, fine-grained, and adaptive access control.

Digital art data transactions differ from ordinary physical transactions. Once data is delivered, the buyer may have copied, cached, parsed, or used it for model training; therefore, failed or rolled-back transactions cannot simply "return the data." Secure delivery and transaction rollback mechanisms must simultaneously address data integrity delivery, trustworthy rights transfer, exception handling, verifiable deletion, and synchronized rights return.

Before the transaction begins, both parties agree on the data object, authorization scope, delivery method, acceptance criteria, payment conditions, liability for breach, and rollback conditions through a smart contract.

The transaction contract is expressed as, as shown in Eq. (25):

$$\text{Contract} = \left(ID_T, ID_d, U_o, U_b, \text{Right}, \text{Policy}, \text{Price}, \text{Deadline}, \text{Rollback}\right) \tag{25}$$

In Eq. (25), IDT denotes the transaction identifier, Right denotes the authorization rights type, Policy denotes usage constraints, Price denotes transaction amount, Deadline denotes the delivery deadline, and Rollback denotes rollback conditions.

The secure delivery process includes four stages: first, the sharer submits the data digest, ownership certificate, and commitment proof; second, the buyer escrows funds in the payment contract; third, the sharer delivers data ciphertext or decryption materials through an off-chain encrypted channel; and fourth, the buyer completes integrity verification and submits a receipt proof. The payment contract releases funds only after verification of the data commitment, ownership certificate, authorization policy, and receipt proof all succeeds.

After receiving the data, the buyer needs to verify that it is consistent with the transaction contract. For complete data delivery, the buyer computes the data digest, as shown in Eq. (26):

$$Hash\left(d'\right) = H\left(d'\right) \tag{26}$$

We compares it with the on-chain registered digest Hashd. If the following condition is satisfied, as shown in Eq. (27):

$$\text{Hash}\left(d'\right) = \text{Hash}_d \tag{27}$$

Then the delivered data is complete and consistent. For digital art data delivered after compression, format conversion, or fragmentation, the system can use a fragment hash tree or perceptual hash to assist verification, thereby reducing misjudgment caused by format changes.

At the same time, the system verifies the consistency among ownership certificate Certd, authorization policy Policyd, and transaction contract Contract. If the data is complete but the authorization scope is mismatched, for example, the buyer only obtains preview rights but receives a commercially usable original file, the system still determines abnormal delivery and triggers review.

Transaction exceptions mainly include four types: the data submitted by the sharer is inconsistent with the commitment; the buyer refuses to submit a receipt or maliciously denies receipt; cross-chain state synchronization fails, causing inconsistent payment states; and ownership disputes or infringement risks are discovered after the transaction is completed.

For the first type of exception, the system determines the sharer's responsibility based on on-chain commitments, zero-knowledge proofs, and data digests, and deducts its pledge. For the second type, the system determines the buyer's responsibility based on delivery logs, receipt proofs, encrypted fragment receipt records, and audit contract status. For the third type, the system suspends settlement and re-verifies multi-chain states through the relay chain. For the fourth type, the system freezes transaction proceeds and subsequent authorizations and enters the ownership arbitration process.

When a transaction is revoked or rolled back, the system must implement fund rollback, rights return, and data deletion. Because digital art data is replicable, the system cannot determine whether the buyer has deleted the data solely based on the buyer's declaration; therefore, a verifiable deletion mechanism is required.

For encrypted delivery data, the system can realize logical deletion by destroying or invalidating the data decryption key Kd and generating a key destruction proof, as shown in Eq. (28):

$$\text{Proof}_{\text{del}} = \text{Sign}_{U_b}\left(ID_T, ID_d, \text{Destroy}\left(K_d\right), T_{\text{del}}\right) \tag{28}$$

For locally cached data, the system requires the buyer to perform deletion and submit remote proof, log proof, or trusted execution environment proof. After being verified by the audit contract, the deletion proof is written on-chain as the basis for transaction rollback. Synchronized rights return is completed through the ownership confirmation system. If the transaction involves ownership transfer, the authority invalidates the new ownership certificate and restores the original ownership certificate state. If the transaction involves copy authorization, the copy authorization certificate is revoked and the revocation record is written on-chain. The rollback record is expressed as, as shown in Eq. (29):

$$\text{RollbackRecord} = \text{Sign}_{CA}\left(ID_T, ID_d, U_o, U_b, \text{Revoke}, \text{Proof}_{\text{del}}, T_r\right) \tag{29}$$

To verify the effectiveness of the secure delivery and transaction rollback mechanism, this paper formally analyzes four dimensions: delivery integrity, payment fairness, responsibility traceability, and rollback verifiability.

(1) Delivery Integrity

Let the data digest registered by the sharer be Hashd, and let the data received by the buyer be d'. If the following condition is satisfied, as shown in Eq. (30):

$$H\left(d^{'}\right)=\text{Hash}_{d}\wedge\text{Verify}\left(\text{Cert}_{d}\right)=1\wedge\text{Verify}\left(CM,\pi_{d}\right)=1 \quad (30)$$

Then the delivered data is consistent with the on-chain registration, and the ownership certificate and commitment proof are both valid. This mechanism supports automatic data integrity verification by the buyer without manual intervention.

(2) Payment Fairness

Let Pay(Ub) denote the buyer's fund escrow operation, let the delivery term denote the sharer's data delivery behavior, and let Receipt(Ub) denote the buyer's submission of a valid receipt proof. The fund release condition of the payment contract satisfies, as shown in Eq. (31):

$$\text{Release(Payment)}=1\Leftrightarrow\text{Pay}\left(U_{b}\right)\wedge\text{Deliver}\left(U_{o},d^{'}\right)\wedge\text{VerifyReceipt}\left(\pi_{r}\right)=1 \quad (31)$$

If data delivery is abnormal or receipt verification fails, Release(Payment)=0. This rule effectively prevents profit from empty delivery and refusal to pay after receipt, thereby protecting the rights and interests of both transaction parties.

(3) Responsibility Traceability

Define the transaction evidence set as, as shown in Eq. (32):

$$E_{T}=\left\{ID_{T},\ \text{Cert}_{d},\ \text{CM},\pi_{d},\pi_{r},\ \text{Policy}_{d},\ \text{Log}_{T},\ \text{Time}_{T}\right\} \quad (32)$$

In Eq. (32), LogT denotes the transaction log and TimeT denotes the timestamp. When a transaction dispute occurs, the arbitrator traces the whole process according to, as shown in Eq. (33):

$$\text{Trace}\left(ID_{T}\right)=\text{Verify}\left(E_{T}\right) \quad (33)$$

Relying on the tamper-resistant property of the chain, the complete evidence set can be verified to reconstruct the whole transaction process, enabling behavior auditing and responsibility determination.

(4) Rollback Verifiability

Let Revoke(Certd) denote certificate revocation, Destroy(Kd) denote decryption key destruction, and Proofdel denote the data deletion proof. The condition for effective transaction rollback is, as shown in Eq. (34):

$$\text{Rollback}\left(ID_{T}\right)=1\Leftrightarrow\text{Revoke}\left(\text{Cert}_{d}\right)\wedge\text{Destroy}\left(K_{d}\right)\wedge\text{Verify}\left(\text{Proof}_{\text{del}}\right)=1 \quad (34)$$

When all conditions hold, the authorization is deemed invalid and transaction rollback is completed; if any condition is not satisfied, rollback is suspended and secondary auditing and arbitration are initiated. In summary, by combining commitment verification, receipt evidence, on-chain traceability, and verifiable deletion, the proposed mechanism realizes complete delivery, fair payment, traceable responsibility, and verifiable rollback, providing formal security support for cross-domain circulation of digital art data.

5. Privacy Protection, Regulatory Arbitration, and Experimental Evaluation

Based on the preceding cross-domain ownership verification, privacy-preserving transactions, and dynamic access control mechanisms, this section further studies layered on-demand privacy protection strategies, low-overhead evidence storage, and violation traceability mechanisms for cross-domain circulation of digital art data. A prototype experimental environment is built to evaluate functional effectiveness, performance overhead, and security. Digital art platform data can be divided into content-type data such as images, videos, and 3D models, and process-type data such as prompts, style parameters, user behaviors, and transaction records. A single privacy protection scheme is difficult to adapt to all scenarios. Therefore, this paper designs layered protection strategies according to data sensitivity levels, access risks, and regulatory requirements, balancing security and data usability.

5.1. Layered On-Demand Privacy Protection Strategy

Digital art data faces differentiated privacy risks at different circulation stages: public works are vulnerable to copyright theft; user behavior data faces profiling leakage risks; prompts and style parameters may reveal commercial creativity; and model-related data concerns core platform assets. Therefore, this paper proposes a layered on-demand privacy protection strategy that dynamically selects corresponding technical solutions according to data sensitivity levels and operation risks.

Let the privacy sensitivity level of digital art data object d be defined as, as shown in Eq. (35):

$$\text{Level}(d)\in\left\{L_{0},L_{1},L_{2},L_{3},L_{4}\right\} \quad (35)$$

In Eq. (35), L0 denotes public data, L1 denotes low-sensitive data, L2 denotes generally sensitive data, L3 denotes highly sensitive data, and L4 denotes core sensitive data.

Based on the preceding cross-domain ownership verification, privacy-preserving transactions, and dynamic access control mechanisms, this section further studies layered on-demand privacy protection, low-overhead evidence storage, and violation traceability for distributed digital art data circulation, and builds a prototype environment to evaluate functional effectiveness, performance overhead, and security. The experiments focus on verifying the applicability of the proposed scheme to privacy protection, cross-chain auditing, dynamic access control, and distributed AI model asset accountability. **This relationship is formalized in Eq. (36).**

$$\text{Re}q=(s,d,o,c) \quad (36)$$

In Eq. (36), s is the access subject, d is the data object, o is the operation type, and c is the access context. The system selects a privacy protection strategy based on data

level, access risk, operation purpose, and compliance constraints, as shown in Eq. (37):

$$P_{priv} = \Phi(\text{Level}(d), R(s,d,o,c), \text{Purpose}(o), \text{Reg}(c)) \tag{37}$$

In Eq. (37), R(s, d, o, c) denotes the comprehensive access risk, Purpose(o) denotes the operation purpose, Reg(c) denotes regulatory constraints, and the mapping function denotes the strategy selection function.

This paper follows the principle of lightweight protection for low risk and strong protection for high risk, avoiding resource waste caused by globally high-intensity encryption. Low-level data mainly uses access control, logs, and watermarking; medium-level data combines encryption, desensitization, and differential privacy; high-level data adopts secret sharing, secure multi-party computation, zero-knowledge proofs, and cross-chain auditing; and core sensitive data further overlays multi-party authorization, threshold recovery, and model watermarking mechanisms.

Privacy protection technologies incur computation, communication, storage, and latency overhead. This paper abstracts strategy selection as a constrained optimization problem that minimizes comprehensive overhead while satisfying the security baseline, as shown in Eq. (38):

$$\begin{aligned} &\min_{\text{Priv}} C\left(P_{\text{priv}}\right) = C_{\text{comp}} + C_{\text{comm}} + C_{\text{store}} + C_{\text{delay}} \\ &\text{s.t.} \quad \text{Risk}(d,o,c) \le \varepsilon_{\text{risk}}, \quad \text{Privacy}\left(P_{\text{priv}}\right) \ge \varepsilon_{\text{priv}} \end{aligned} \tag{38}$$

In Eq. (38), Ccomp, Ccomm, Cstore, and Cdelay denote computation, communication, storage, and latency overhead, respectively; ε_{risk} denotes the acceptable risk threshold, and ε_{priv} denotes the minimum privacy protection strength threshold. For real-time interactive services such as image generation and online editing, low-latency protection schemes are preferred. For copyright registration, high-value transactions, and regulatory auditing, more secure technologies such as zero-knowledge proofs, secret sharing, and cross-chain auditing are adopted to dynamically adapt privacy protection strategies to business scenarios.

5.2. Low-Overhead Evidence Storage and Violation Traceability

Directly storing raw digital art files, models, and full transaction data on-chain would lead to high storage costs, privacy leakage risks, and low on-chain processing efficiency. Therefore, this paper designs a low-overhead evidence storage and violation traceability mechanism that combines off-chain storage, on-chain digests, evidence indexes, and watermark-based tracing to realize low-cost, verifiable, and accountable cross-domain data governance.

For data object d, the system stores only a lightweight digest on-chain. The data evidence digest is defined as, as shown in Eq. (39):

$$E_d = \left(ID_d, H(d), H\left(M_d\right), \text{Cert}_d, \text{Policy}_d, T_d, \text{Sig}_o\right) \tag{39}$$

In Eq. (39), IDd denotes the data asset identifier, H(d) denotes the data content hash, H(Md) denotes the metadata hash, Certd denotes the ownership certificate digest, Policyd denotes the policy digest, Td denotes the timestamp, and Sigo denotes the owner's signature. For a single cross-domain transaction T, transaction evidence is generated as, as shown in Eq. (40):

$$E_T = \left(ID_T, ID_d, U_o, U_b, CM, \pi_d, \pi_r, \text{Policy}_T, \text{State}_T, T_T\right) \tag{40}$$

In Eq. (40), ID_T denotes the transaction identifier, CM denotes the data commitment, π_d denotes the authenticity proof, π_r denotes the receipt proof, Policy_T denotes the transaction policy, StateT denotes the transaction state, and T_T denotes the transaction timestamp.

This scheme adopts a lightweight model of storing digests on-chain, ciphertext off-chain, and verifying the original text on demand. Only verifiable digests and indexes are retained on-chain, avoiding exposure of raw data. Encrypted data, complete metadata, and detailed logs are uniformly stored off-chain. When a dispute or regulatory review occurs, off-chain data can be retrieved based on on-chain indexes to complete consistency verification.

Data circulation includes key stages such as generation, registration, authorization, access, transaction, publication, and auditing. The system generates an evidence unit for each event, as shown in Eq. (41):

$$e_k = H\left(ID_d \| \text{Event}_k \| \text{Actor}_k \| \text{Time}_k \| \text{State}_k\right) \tag{41}$$

In Eq. (41), Eventk denotes the event type, Actork denotes the participating subject, Timek denotes the occurrence time, and Statek denotes the event state. All events are organized in chronological order to form an evidence chain, as shown in Eq. (42):

$$\text{Chain}_d = \{e_1, e_2, \ldots, e_n\} \tag{42}$$

To prevent content tampering, event deletion, and sequence tampering, a hash-chain structure is further adopted, as shown in Eq. (43):

$$e_k = H\left(e_{k-1} \| ID_d \| \text{Event}_k \| \text{Actor}_k \| \text{Time}_k \| \text{State}_k\right) \tag{43}$$

If any link is tampered with, all subsequent hash units fail verification, thereby guaranteeing the integrity and traceability of the circulation process.

This paper targets violations such as unauthorized access, illegal downloading, out-of-scope use, copyright infringement, model abuse, transaction repudiation, and unauthorized dissemination, and performs joint traceability by integrating evidence chains, access logs, on-chain evidence, and watermark information. For suspicious data d', watermarks, perceptual hashes, semantic features, and dissemination metadata are extracted to form a traceability feature set, as shown in Eq. (44):

$$\mathrm{Trace}\left(d^{'}\right)=\left(W_{d^{'}},PH_{d^{'}},S_{d^{'}},\mathrm{Meta}_{d^{'}}\right) \tag{44}$$

In Eq. (44), the corresponding terms denote watermark information, perceptual hash, semantic features, and dissemination metadata, respectively. These features are then matched against on-chain registered data:

In the prototype, watermark evidence is generated in two complementary ways. For digital artworks, a robust imperceptible watermark is embedded into mid-frequency image components and detected by correlation matching after compression, cropping, or format conversion. For AIGC model services, a model-level fingerprint is associated with generated outputs through watermark identifiers and prompt/session digests. The detector returns a confidence score in [0,1], which is fused with perceptual hash similarity and semantic similarity. If the fused score is above τtrace, the evidence chain is traversed to determine whether the detected copy is within the authorized purpose, time range, and dissemination path. This relationship is formalized in Eq. (45).

$$\mathrm{Match}\left(d^{'},d\right)=\omega_1\,\mathrm{Sim}\left(W_{d^{'}},W_d\right)+\omega_2\,\mathrm{Sim}\left(PH_{d^{'}},PH_d\right)+\omega_3\,\mathrm{Sim}\left(S_{d^{'}},S_d\right)+\omega_4\,\mathrm{Sim}\left(\mathrm{Meta}_{d^{'}},\mathrm{Meta}_d\right) \tag{45}$$

In Eq. (45), the corresponding term denotes the weight coefficient. When the matching score reaches the threshold, the suspicious data is considered associated with the original data. The evidence chain Chaind is then used to trace back legal authorization, transaction, and dissemination paths to locate the responsible subject.

When transaction disputes, copyright disputes, or privacy leakage events occur, regulators and arbitrators can rely on on-chain evidence, off-chain operational logs, and watermark detection results to make compliance decisions. The process includes evidence submission, validity verification, responsibility localization, and violation handling. First, the disputing party submits transaction identifiers, data identifiers, ownership certificates, suspicious works, and access logs. Then, the system verifies on-chain data digests, digital signatures, transaction commitment proofs, receipt credentials, and timestamps. Next, complete data evidence chains and violation traceability results are combined to accurately locate the violation node, violation type, and responsible subject. Finally, the system executes authorization revocation, fund freezing, violation compensation, blacklist updating, certificate invalidation, and regulatory alerts according to predefined smart contract rules. This arbitration mechanism can complete evidence verification and responsibility determination without disclosing raw sensitive data, effectively balancing privacy protection and regulatory arbitration requirements in cross-domain digital art data transactions.

5.3. Experimental Environment and Evaluation Metrics

To verify the feasibility and effectiveness of the proposed mechanism, this section builds a prototype system and designs a comprehensive experimental evaluation scheme. Rather than focusing on the performance comparison of a single cryptographic algorithm, the experiments are based on realistic cross-domain data circulation scenarios in digital art platforms and comprehensively evaluate the overall effectiveness and practicality of the proposed framework from multiple dimensions, including functional correctness, system overhead, dynamic access control performance, privacy protection capability, and violation traceability.

The experimental prototype system is divided into three logically isolated business domains: the digital art design platform domain, the AIGC model service domain, and the copyright transaction and regulatory domain. Each domain independently deploys consortium-chain nodes, and the relay chain completes cross-domain transaction state synchronization, data interaction, and cross-chain audit verification. The system includes data providers, data buyers, model service providers, copyright agencies, regulators, and other business subjects, fitting real cross-domain circulation scenarios.

The prototype deployment environment is as follows: the consortium chain uses the FISCO BCOS architecture to implement cross-domain evidence storage, ownership registration, and smart contract deployment and execution; the relay-chain module is responsible for cross-chain data query, transaction state synchronization, and audit contract invocation; the off-chain storage system encrypts and stores digital artworks, AIGC model files, operation logs, and metadata; the privacy protection module integrates hash commitments, secret sharing, zero-knowledge proofs, and multidimensional watermark-based tracing; and the access control module implements adaptive dynamic authorization decisions based on subject reputation, data sensitivity level, operation risk, and contextual environment.

The prototype was deployed on a workstation with an Intel i7-class CPU, 32 GB RAM, Ubuntu 22.04, Python 3.10, Solidity-compatible smart-contract scripts, and three logically isolated consortium-chain domains. Each domain contains four consortium-chain nodes and one off-chain storage node. The relay module maintains cross-chain transaction state indexes and synchronizes state every 200 ms in the default setting. Hash commitments use SHA-256 digests and Pedersen-style commitment simulation over a 256-bit prime-order group; Shamir secret sharing uses n=5 and t=3; zero-knowledge proof generation and verification are simulated with proof-size and latency parameters calibrated from typical ZKP libraries; watermark matching uses normalized correlation and perceptual-hash similarity; and all reported values are averaged over five repeated runs.

The experimental data includes two categories. The first consists of public illustrations, art images, style-transfer works, and AIGC-generated images, which are used to simulate work registration, authorized transactions, and infringement traceability. The second

consists of artificially synthesized transaction logs, user access records, authorization policies, and subject reputation data, which are used to verify cross-domain access control, privacy-preserving transactions, and regulatory arbitration processes. All raw data involving user privacy is anonymized to avoid privacy leakage and ensure experimental compliance.

To support statistical validation, the test set contains 5,000 registered digital art objects, including 2,000 public illustration images, 1,000 style-transfer images, 1,000 AIGC-generated images, 500 prompt/metadata records, and 500 model-service invocation records. Access-control experiments contain 20,000 requests with legal, low-risk, high-frequency, and malicious behavior labels. Transaction experiments contain 3,000 cross-domain transactions. Violation-traceability experiments contain 400 suspicious dissemination samples. Each baseline uses the same data split and transaction workload to ensure fair comparison.

To comprehensively verify the performance advantages of the proposed framework, three mainstream schemes are set as baselines for horizontal comparison with the proposed scheme.

Baseline-1 (traditional centralized transaction scheme): A centralized database is used to uniformly record transaction information and authorization states. Blockchain evidence storage, cross-chain auditing, zero-knowledge proofs, and other security mechanisms are not introduced, and transaction governance relies on a centralized third party.

Baseline-2 (ordinary blockchain evidence storage scheme): Only data hash values and transaction records are stored on-chain. A single-chain architecture provides basic evidence storage, without support for cross-chain state synchronization, privacy-preserving transaction auditing, or abnormal transaction rollback.

Baseline-3 (static access control scheme): Access authorization is completed based on fixed roles and static permission rules, without considering dynamic changes in subject reputation, differences in data sensitivity levels, or contextual access risks, and the permission policy cannot adaptively adjust.

Baseline-5 (privacy-preserving transaction scheme): this baseline uses commitments and encrypted delivery to protect transaction privacy, but it does not integrate DID/VC identity verification, dynamic trust updating, watermark-based traceability, or regulatory arbitration.

Baseline-4 (blockchain copyright management scheme): this baseline registers copyright certificates and data hashes on a single chain and supports ownership query and infringement evidence retrieval, but does not support privacy-preserving transaction proofs, cross-chain state synchronization, dynamic access control, or rollback verification.

Proposed (the scheme in this paper): The proposed scheme integrates layered on-demand privacy protection, risk-aware dynamic access control, cross-consortium-chain privacy-preserving transaction auditing, low-overhead on-chain evidence storage, violation traceability, and smart-contract-based regulatory arbitration, realizing full-process security governance for cross-domain circulation of digital art data.

Through comparisons among multiple schemes, the advantages of the proposed scheme in cross-domain trusted interaction, privacy security protection, dynamic precise authorization, transaction auditability, and violation traceability accountability can be effectively verified.

This paper constructs an evaluation metric system from three dimensions: functionality, performance, and security, to comprehensively quantify the overall performance of the framework.

(1) Functional verification metrics: ownership confirmation success rate, transaction audit pass rate, authorization decision correctness, violation traceability success rate, and transaction rollback success rate. The ownership confirmation success rate measures the ability of the system to identify and authenticate ownership of original works, edited works, and AIGC-derived works. The violation traceability success rate evaluates the ability of the system to accurately locate responsible subjects based on watermarks, perceptual hashes, semantic features, and evidence chains.

(2) Performance overhead metrics: ownership confirmation time, transaction audit latency, access decision latency, on-chain evidence storage overhead, off-chain storage overhead, communication overhead, and smart contract execution overhead. For real-time interaction in digital art platforms, access decision latency and model invocation latency are core metrics for user experience and system practicality.

(3) Security metrics: privacy leakage risk coefficient, unauthorized access interception rate, transaction repudiation resistance, evidence tampering detection rate, copyright infringement identification rate, and transaction rollback verifiability. These metrics verify the ability of the framework to resist malicious attacks such as identity forgery, data fraud, payment repudiation, illegal dissemination, and evidence tampering.

The core evaluation metrics of this paper are formally defined as follows. Access decision accuracy, as shown in Eq. (46):

$$Acc_{ac} = \frac{TP + TN}{TP + TN + FP + FN} \tag{46}$$

In Eq. (46), TP denotes the number of legal accesses correctly allowed, TN denotes the number of illegal accesses correctly intercepted, FP denotes the number of illegal accesses incorrectly allowed, and FN denotes the number of legal accesses incorrectly denied. Violation traceability success rate, as shown in Eq. (47):

$$SR_{\text{trace}} = \frac{N_{\text{correct}}}{N_{\text{violation}}} \tag{47}$$

In Eq. (47), Ncorrect denotes the number of violation events for which the responsible subject is successfully located, and Nviolation denotes the total number of violation events.

On-chain evidence storage compression ratio, as shown in Eq. (48):

$$CR_{\text{chain}} = 1 - \frac{\text{Size}(E_{\text{chain}})}{\text{Size}(D_{\text{raw}})} \quad (48)$$

In Eq. (48), Size(Echain) denotes the size of lightweight on-chain evidence, and Size(Draw) denotes the size of raw art data. This metric measures the storage reduction effect of the evidence storage scheme.

For statistical reliability, each latency and throughput metric is reported as the mean value of five repeated runs, and the standard deviation is recorded when the metric is sensitive to concurrency. For classification-like metrics such as access decision accuracy and traceability success rate, precision, recall, and F1 score are also computed where applicable. This relationship is formalized in Eq. (49).

$$T_{\text{audit}} = T_{\text{commit}} + T_{\text{proof}} + T_{\text{verify}} + T_{\text{cross}} + T_{\text{contract}} \quad (49)$$

In Eq. (49), Tcommit denotes the time for generating a data commitment, Tproof denotes the time for generating a zero-knowledge proof, Tverify denotes the time for evidence verification, Tcross denotes the time for cross-chain synchronization, and Tcontract denotes the time for smart contract execution.

5.4. Functional Verification, Performance Overhead, and Security Analysis

To verify the effectiveness of the proposed mechanisms, this paper constructs a prototype simulation environment for cross-domain circulation of digital art data. The experiments cover digital art work ownership confirmation, cross-consortium-chain privacy-preserving transaction auditing, dynamic access control, low-overhead evidence storage, and violation traceability. The experiments simulate 5,000 digital art data objects, 20,000 cross-domain access requests, 3,000 cross-domain authorization transactions, and 400 violation dissemination samples. The data objects include original images, style-transfer images, AIGC-generated images, prompt digests, copyright credentials, and transaction records. The baseline schemes include a centralized transaction scheme, an ordinary blockchain evidence storage scheme, a static access control scheme, and the proposed scheme.

The experimental evaluation focuses on cross-domain trusted interaction, privacy protection overhead, cross-chain state synchronization, dynamic authorization effects, and evidence traceability in distributed scenarios, so as to verify the applicability of the proposed mechanisms to data security problems in emerging distributed networks and systems.

First, this paper verifies the arbitrable damage-tolerant ownership verification method. The original digital artworks are subjected to compression, cropping, format conversion, partial occlusion, style transfer, and combined perturbations, and the system is tested to determine whether it can identify the ownership association between the perturbed work and the original registered work. The experimental results are shown in Table 2.

Table 2. Identification Results of Art Data Right Confirmation Under Different Disturbance Conditions

Disturbance Type	Number of Test Samples	Right Confirmation Success Rate /%	Average Similarity	Misjudgment Rate /%
No Disturbance	500	100.0	0.982	0.0
Image Compression	500	98.7	0.941	1.3
Format Conversion	500	97.9	0.932	2.1
Partial Cropping	500	94.5	0.891	5.5
Partial Occlusion	500	91.2	0.853	8.8
Style Transfer	500	88.6	0.827	11.4
Combined Disturbance	500	85.9	0.796	14.1

As shown in Table 2, the proposed method still maintains a high ownership confirmation success rate under mild perturbations such as image compression and format conversion. Under partial occlusion, style transfer, and combined perturbations, the success rate decreases, but the method still maintains good damage-tolerant identification capability through joint matching of content features, semantic features, watermark information, and on-chain certificates. This shows that the proposed method is more suitable than traditional single-hash ownership verification for derivative creation, format changes, and cross-platform dissemination of digital artworks. Second, the privacy-preserving transaction audit mechanism across consortium chains is verified. Five scenarios are tested: normal transaction, data tampering, invalid proof, missing receipt, and payment repudiation. The results are shown in Table 3.

Table 3. Verification Results of Privacy Transaction Audit Function Across Consortium Chains

Transaction Scenario	Number of Test Transactions	Audit Pass Rate /%	Anomaly Detection Rate /%	Fund Mis-release Rate /%
Normal Transaction	1000	99.3	—	0.0
Data Tampering	500	0.0	100.0	0.0
Invalid Zero-Knowledge Proof	500	0.0	100.0	0.0

Receipt Missing	500	0.0	100.0	0.0
Payment Repudiation	500	0.0	98.8	0.0

As shown in Table 3, the proposed scheme can complete data authenticity verification, delivery integrity verification, and receipt verification without disclosing raw art data. When data is tampered with or the proof is invalid, the audit contract can reject transaction settlement. When the buyer attempts payment repudiation, the system can determine responsibility based on receipt proofs, on-chain states, and delivery logs. The fund mis-release rate remains 0.0% in all anomalous scenarios, confirming that the cross-chain privacy-preserving transaction audit mechanism effectively ensures fairness and non-repudiation in cross-domain transactions. Third, the risk-aware dynamic access control mechanism is verified. The experiment constructs four types of subjects—legal users, low-risk users, high-frequency access users, and malicious users—and sets five operation types: preview, download, commercial publication, model training, and reauthorization. The authorization effects of static access control and the proposed dynamic access control scheme are compared, and the results are shown in Table 4.

Table 4. Verification Results of Privacy Transaction Audit Function Across Consortium Chains

Transaction Scenario	Number of Test Transactions	Audit Pass Rate /%	Anomaly Detection Rate /%	Fund Mis-release Rate /%
Normal Transaction	1000	99.3	—	0.0
Data Tampering	500	0.0	100.0	0.0
Invalid Zero-Knowledge Proof	500	0.0	100.0	0.0
Receipt Missing	500	0.0	100.0	0.0
Payment Repudiation	500	0.0	98.8	0.0

Table 5. Authorization Decision Effect of Different Access Control Schemes

Scheme	Decision Accuracy Rate /%	Illegal Access Interception Rate /%	False Rejection Rate of Legitimate Access /%	F1 Score
Static RBAC	81.7	76.4	13.8	0.803
Static ABAC	86.5	82.1	10.6	0.861
Risk-aware Access Control	92.8	91.3	6.4	0.927
Proposed Mechanism in This Paper	95.6	94.9	4.8	0.954

The results in Table 4 show that the proposed scheme can complete data authenticity verification, delivery integrity verification, and receipt verification without disclosing raw art data. When data is tampered with or the proof is invalid, the audit contract can reject transaction settlement. When the buyer attempts payment repudiation, the system can determine responsibility based on receipt proofs, on-chain states, and delivery logs. Therefore, the proposed scheme can effectively ensure fairness and non-repudiation in cross-domain transactions. Third, the risk-aware dynamic access control mechanism is verified. The experiment constructs four types of subjects—legal users, low-risk users, high-frequency access users, and malicious users—and sets five operation types: preview, download, commercial publication, model training, and reauthorization. The authorization effects of static access control and the proposed dynamic access control scheme are compared, and the results are shown in Table 5.

As shown in Table 5, the proposed mechanism outperforms static RBAC and static ABAC in decision accuracy, illegal access interception rate, and F1 score. The main reason is that the proposed mechanism considers not only subject roles and static attributes but also subject reputation, historical behavior, data sensitivity level, operation risk, and contextual state, enabling more precise identification of high-risk access behaviors. Compared with ordinary risk-aware access control, the proposed mechanism further combines access feedback and trust updating, allowing the system to dynamically adjust authorization results according to user behavior.

To evaluate the system overhead of the proposed scheme, this paper tests ownership confirmation time, access decision latency, cross-chain audit latency, and on-chain storage overhead. First, ownership confirmation processing time under different sizes of digital art data is tested. The results are shown in Table 6.

Table 6. Ownership Confirmation Processing Time Under Different Data Sizes

Data Size / MB	Feature Extraction / ms	Arbitration Code Generation / ms	Certificate Issuance / ms	On-chain Storage / ms	Total Time / ms
1	38.4	4.7	12.3	96.5	151.9
5	91.6	5.1	12.6	98.2	207.5
10	164.8	5.4	12.8	101.7	284.7
20	301.5	6.0	13.1	106.4	427.0
50	724.9	7.2	13.5	112.8	858.4

As shown in Table 6, the total ownership confirmation time increases with data size, and the main overhead comes from art image feature extraction. Arbitration code generation and certificate issuance mainly involve hash computation and digital signatures, and their time overhead is small and does not vary significantly with data size. The on-chain storage time remains relatively stable, indicating that the proposed on-chain digest evidence storage method can avoid the performance bottleneck caused by directly storing large-scale data on-chain. Next, access control decision latency is tested. The results are shown in Table 7.

Table 7. Average Decision Latency of Different Access Control Schemes

Scheme	Identity Verification / ms	Policy Matching / ms	Risk Assessment / ms	Trust Update / ms	Total Latency / ms
Static RBAC	4.2	2.1	—	—	6.3
Static ABAC	5.6	6.8	—	—	12.4
Risk-aware Access Control	7.8	7.1	8.4	—	23.3
Proposed Mechanism in This Paper	8.1	7.4	8.6	3.2	27.3

Table 8. Latency Breakdown of Cross-Consortium Chain Private Transaction Auditing

Audit Stage	Average Latency / ms
Data Commitment Generation	9.6
Zero-Knowledge Proof Generation	82.4
Zero-Knowledge Proof Verification	31.5
Secret Sharing Processing	46.7
Cross-Chain State Synchronization	326.8
Smart Contract Execution	74.2
Receipt Verification	28.9
Total Auditing Latency	**600.1**

Table 7 shows that the proposed mechanism introduces additional risk assessment and trust updating overhead compared with static access control, but the average decision latency remains at the millisecond level and can meet the online access control requirements of digital art platforms. Considering that the proposed mechanism significantly improves the illegal access interception rate and authorization decision accuracy, the additional overhead is acceptable. The cross-consortium-chain privacy-preserving transaction audit overhead is then tested. The results are shown in Table 8.

As shown in Table 8, cross-chain state synchronization is the main source of transaction audit latency, followed by zero-knowledge proof generation and smart contract execution. Since digital art data transactions are usually not high-frequency millisecond-level transactions but emphasize transaction trustworthiness, copyright protection, and non-repudiation, the average audit latency of approximately 600 ms is acceptable for high-value digital asset transaction scenarios. Cross-chain message caching, batch verification, and proof pre-generation can further reduce audit latency in future work. Finally, the low-overhead evidence storage mechanism is evaluated. The experiment compares the storage overhead of directly storing raw data on-chain with that of the proposed on-chain digest evidence storage method. The results are shown in Table 9.

Table 9. Cost Comparison of On-chain Evidence Storage

Original Data Size / MB	Direct On-chain Storage / KB	On-chain Evidence of Proposed Scheme / KB	Storage Compression Ratio / %
1	1 024	2.1	99.79
5	5 120	2.4	99.95
10	10 240	2.8	99.97
20	20 480	3.2	99.98
50	51 200	3.7	99.99

Table 9 shows that the proposed method, which combines on-chain digests with off-chain ciphertext, can significantly reduce on-chain storage pressure. As the raw data size increases, the on-chain evidence size of the proposed scheme increases only slowly, whereas the storage overhead of direct on-chain storage increases linearly. Therefore, the proposed scheme is more suitable for cross-domain circulation of large-scale digital art data such as images, videos, and model files.

To verify violation traceability capability, the experiments simulate six types of violations: illegal copying, local modification, compressed dissemination, de-watermarking attacks, continued use after authorization expiration, and model training abuse. Single watermark detection, single hash matching, perceptual hash matching, and the proposed multi-evidence fusion traceability mechanism are compared. The results are shown in Table 10.

Table 10. Violation Location Performance of Different Traceability Methods

Traceability Method	Traceability Success Rate / %	False Positive Rate / %	Average Traceability Time / ms
Single Hash Matching	61.4	3.2	18.6
Single Watermark Detection	82.1	5.7	46.3

Perceptual Hash Matching	86.9	7.4	39.8
Watermark + Perceptual Hash	91.5	5.1	64.7
Proposed Multi-evidence Fusion Scheme	94.8	3.9	78.5

The results in Table 10 show that single hash matching is very sensitive to minor data modifications and cannot cope with common dissemination changes such as compression, cropping, and format conversion. Single watermark detection suffers from degraded recognition capability under de-watermarking attacks. Perceptual hashing has certain robustness to slight visual changes but a relatively high false positive rate. The proposed multi-evidence fusion mechanism combines watermarks, perceptual hashes, semantic features, access logs, and on-chain evidence, achieving the highest traceability success rate while maintaining a low false positive rate. Furthermore, the evidence verification success rate of the regulatory arbitration process is tested under different dispute types. The results are shown in Table 11.

Table 11. Evidence Verification Results for Regulatory Arbitration

Dispute Type	Number of Test Samples	Evidence Verification Success Rate / %	Responsible Party Localization Rate / %
Ownership Dispute	100	96.0	94.0
Transaction Repudiation	100	99.0	98.0
Unauthorized Access	100	95.0	93.0
Illegal Dissemination	100	94.0	92.0
Model Misuse	100	91.0	88.0

As shown in Table 9, the proposed mechanism performs best in transaction repudiation scenarios because commitments, proofs, receipts, and contract states in the transaction process provide clear on-chain evidence. In model abuse scenarios, the responsible party localization rate is relatively lower, mainly because model training or inference abuse is difficult to observe directly and must be inferred from invocation logs, model watermarks, and output content features. This result indicates that misuse detection for AIGC model assets remains a key direction for future improvement.

The formal security properties can therefore be summarized as follows. Privacy preservation holds because audit verification uses commitments and proofs instead of plaintext data; non-repudiation holds because payment release requires valid receipt proof and contract state consistency; ownership robustness holds because the ownership decision uses multi-evidence similarity rather than a single fragile hash; rollback verifiability holds when certificate revocation, key destruction, deletion proof, and rollback records are simultaneously valid; and collusion resistance holds as long as fewer than t secret shares are compromised and at least one audit node in each transaction path remains honest. These properties do not eliminate all risks, but they make violations detectable, attributable, and arbitrable.

The security analysis is organized under the adversarial model defined in Section 3.1.2. Against A1, DID/VC verification, certificate validity checking, and signed evidence records prevent forged identities from being accepted unless the attacker compromises valid credentials. Against A2, raw prompts, artworks, model inputs, and decryption keys are not submitted to the audit chain in plaintext; commitments, key sharing, encrypted delivery, and zero-knowledge proof verification reduce the information available to semi-honest service providers. Against A3, receipt proofs, delivery logs, watermark identifiers, and evidence-chain timestamps provide non-repudiation evidence, making payment denial and unauthorized dissemination traceable. Against A4, commitment binding, data digest verification, and certificate-policy consistency checks prevent dishonest sharers from replacing registered data with forged content. Against A5, relay-chain state indexes and contract state re-verification reduce the risk of inconsistent settlement when cross-chain synchronization is delayed.

The experimental results show that the proposed scheme improves the security of cross-domain digital art data circulation from five aspects: identity, data, transaction, evidence, and copyright.

First, in terms of identity security, the DID and verifiable credential mechanism can prevent subjects with invalid identities, expired credentials, or unsatisfied attribute conditions from initiating high-risk access. Combined with dynamic risk assessment, the system achieves an illegal access interception rate of 94.9%, which is significantly better than the static access control scheme.

Second, in terms of data authenticity, Pedersen commitments and zero-knowledge proofs ensure that transaction data is consistent with on-chain commitments. When data is tampered with or the proof is invalid, the audit contract can detect anomalies and reject settlement, resulting in a fund mis-release rate of 0.

Third, in terms of payment fairness, smart contracts use fund escrow, receipt verification, and conditional release mechanisms, requiring the sharer to complete valid delivery before receiving payment and making it difficult for the buyer to deny the transaction after receiving the data.

Fourth, in terms of evidence integrity, on-chain digests, hash chains, and timestamps can detect evidence tampering. If any transaction event is modified, subsequent evidence chain verification will fail, thereby

ensuring the credibility of evidence in regulatory arbitration.

Fifth, in terms of copyright accountability, the multi-evidence fusion traceability mechanism jointly matches watermarks, perceptual hashes, semantic features, access logs, and on-chain evidence, enabling a high success rate in responsibility localization for illegal dissemination and authorization violations.

The comprehensive experimental results show that the proposed scheme has the following advantages compared with the baseline schemes.

Compared with the centralized transaction scheme, the proposed scheme introduces on-chain evidence storage, cross-chain auditing, and privacy proofs. Although it incurs certain computation and communication overhead, it significantly improves transaction verifiability and non-repudiation. The centralized scheme is efficient for low-risk transactions, but it has obvious shortcomings in cross-platform dispute arbitration, evidence trustworthiness, and copyright accountability.

Compared with the ordinary blockchain evidence storage scheme, the proposed scheme does not directly store raw art data on-chain; instead, it combines on-chain digests, off-chain ciphertext, and evidence indexes, thereby significantly reducing on-chain storage overhead while ensuring verifiability. Experimental results show that for 50 MB data, the on-chain storage compression ratio of the proposed scheme reaches 99.99%, making it more suitable for large-scale digital art data and model files.

Compared with static access control schemes, the proposed scheme can dynamically authorize access according to subject trustworthiness, data sensitivity level, operation risk, and contextual state. Experimental results show that its authorization decision accuracy reaches 95.6% and its illegal access interception rate reaches 94.9%, which are significantly higher than those of static RBAC and static ABAC.

Compared with single violation traceability methods, the proposed multi-evidence fusion mechanism maintains a high traceability success rate even when works are compressed, cropped, format-converted, or locally tampered with. Experimental results show that the traceability success rate of the proposed mechanism reaches 94.8%, outperforming single hash matching, watermark detection, and perceptual hash matching methods.

However, the proposed scheme still has certain limitations. First, zero-knowledge proof generation, cross-chain state synchronization, and smart contract execution introduce additional audit latency, which may require further optimization for ultra-high-frequency small-value transaction scenarios. Second, the responsible party localization rate for model abuse is still lower than that for ordinary artwork infringement, indicating that watermark embedding, invocation tracing, and output attribution for AIGC model assets should be further strengthened. Finally, the experiments are mainly based on a prototype simulation environment; future work should further verify system stability and scalability using real digital art platform datasets and large-scale user access scenarios.

In summary, the experimental results show that the proposed cross-domain data flow security governance mechanism can realize trustworthy ownership confirmation, privacy-preserving transactions, dynamic access control, low-overhead evidence storage, violation traceability, and regulatory arbitration for digital art data while maintaining acceptable performance overhead, thereby providing effective support for secure data element circulation in digital art platforms.

6. Conclusion

This paper addresses cross-domain data circulation in distributed digital art platforms and proposes an integrated security governance and privacy protection mechanism. In the revised version, the technical novelty is clarified as an application-oriented but executable integration of multi-evidence ownership verification, privacy-preserving cross-consortium-chain transaction auditing, game-theoretic dynamic access control, contract-constrained secure delivery and rollback, low-overhead evidence storage, and watermark-based violation traceability. Compared with existing single-chain copyright evidence storage, static access control, or isolated privacy-preserving transaction methods, the proposed mechanism binds identity, data, policy, commitment, receipt, watermark, and evidence-chain states into a closed-loop lifecycle governance process. Prototype experiments demonstrate that the mechanism achieves robust ownership confirmation under multiple perturbations, improves access decision accuracy to 95.6%, intercepts 94.9% of illegal access, detects major abnormal transaction scenarios with 98.8%–100% effectiveness, and reduces on-chain storage overhead by storing only digest-level evidence. Additional scalability tests show that the access decision module maintains low latency under high concurrency, while the cross-chain transaction audit module remains feasible for medium-scale digital asset transactions. The security analysis further explains privacy preservation, non-repudiation, ownership robustness, rollback verifiability, and bounded collusion resistance under explicit adversarial assumptions. Future work will focus on real-platform deployment, stronger verifiable deletion, batch zero-knowledge proof optimization, heterogeneous-chain interoperability, and privacy-risk estimation for repeated AIGC model interactions.

References

[1] Azcoitia, S. A., & Laoutaris, N. (2022). A survey of data marketplaces and their business models. ACM SIGMOD Record, 51(3), 18–29.

[2] Fernandez, R. C., Subramaniam, P., & Franklin, M. J. (2020). Data market platforms: Trading data assets to solve data problems. Proceedings of the VLDB Endowment, 13(12), 1933–1947.
[3] Zhang, M., Beltran, F., & Liu, J. (2023). A survey of data pricing for data marketplaces. IEEE Transactions on Big Data, 9(4), 1038–1056.
[4] Spiekermann, S., Acquisti, A., Böhme, R., & Hui, K. L. (2015). The challenges of personal data markets and privacy. Electronic Markets, 25(2), 161–167.
[5] Jing, Y., Yang, Y., Feng, Z., Ye, J., Yu, Y., & Song, M. (2020). Neural style transfer: A review. IEEE Transactions on Visualization and Computer Graphics, 26(11), 3365–3385.
[6] Chen, F. L., Zhang, D. Z., Han, M. L., Chen, X. Y., Shi, J., Xu, S., & Xu, B. (2023). VLP: A survey on vision-language pre-training. Machine Intelligence Research, 20(1), 38–56.
[7] Yang, L., Zhang, Z., Song, Y., Hong, S., Xu, R., Zhao, Y., Shao, Y., Zhang, W., Cui, B., & Yang, M. H. (2023). Diffusion models: A comprehensive survey of methods and applications. ACM Computing Surveys, 56(4), 1–39.
[8] Croitoru, F. A., Hondru, V., Ionescu, R. T., & Shah, M. (2023). Diffusion models in vision: A survey. IEEE Transactions on Pattern Analysis and Machine Intelligence, 45(9), 10850–10868.
[9] Wang, S., Zhang, Y., & Zhang, Y. (2018). A blockchain-based framework for data sharing with fine-grained access control in decentralized storage systems. IEEE Access, 6, 38437–38450.
[10] Liu, D., Huang, C., Ni, J., Lin, X., & Shen, X. S. (2022). Blockchain-cloud transparent data marketing: Consortium management and fairness. IEEE Transactions on Computers, 71(12), 3322–3335.
[11] de Vos, M., Ileri, C. U., & Pouwelse, J. (2021). XChange: A blockchain-based mechanism for generic asset trading in resource-constrained environments. World Wide Web, 24(5), 1691–1728.
[12] Mühle, A., Grüner, A., Gayvoronskaya, T., & Meinel, C. (2018). A survey on essential components of a self-sovereign identity. Computer Science Review, 30, 80–86.
[13] Ferdous, M. S., Chowdhury, F., & Alassafi, M. O. (2019). In search of self-sovereign identity leveraging blockchain technology. IEEE Access, 7, 103059–103079.
[14] Li, H. (2026). A cloud environment security access control scheme based on federated learning and fuzzy logic integration. EAI Endorsed Transactions on Scalable Information Systems, 12(9). https://doi.org/10.4108/eetsis.11731
[15] Hu, V. C., Kuhn, D. R., Ferraiolo, D. F., & Voas, J. (2015). Attribute-based access control. Computer, 48(2), 85–88.
[16] Park, J., & Sandhu, R. (2004). The UCONABCusage control model. ACM Transactions on Information and System Security, 7(1), 128–174.
[17] Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. Foundations and Trends in Theoretical Computer Science, 9(3–4), 211–417.
[18] Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ... & Sethi, T. (2021). Advances and open problems in federated learning. Foundations and Trends in Machine Learning, 14(1–2), 1–210.
[19] Guo, Z., Sun, Y., Zhang, X., & Wu, L. (2026). An efficient privacy-preserving secure aggregation scheme for federated learning with input verification and dropout resistance. EAI Endorsed Transactions on Scalable Information Systems, 12(8), 1–14.
[20] Bogdanov, D., Niitsoo, M., Toft, T., & Willemson, J. (2012). High-performance secure multi-party computation for data mining applications. International Journal of Information Security, 11(6), 403–418.
[21] Evans, D., Kolesnikov, V., & Rosulek, M. (2018). A pragmatic introduction to secure multi-party computation. Foundations and Trends in Privacy and Security, 2(2–3), 70–246.
[22] Xue, J., & Yan, W. (2026). Edge computing communication privacy protection method based on federated learning algorithm. EAI Endorsed Transactions on Scalable Information Systems, 12(9). https://doi.org/10.4108/eetsis.12254
[23] Shamir, A. (1979). How to share a secret. Communications of the ACM, 22(11), 612–613.
[24] Abe, M., Fuchsbauer, G., Groth, J., Haralambiev, K., & Ohkubo, M. (2016). Structure-preserving signatures and commitments to group elements. Journal of Cryptology, 29(2), 363–421.
[25] Goldreich, O., & Oren, Y. (1994). Definitions and properties of zero-knowledge proof systems. Journal of Cryptology, 7(1), 1–32.
[26] Xue, M., Zhang, Y., Wang, J., & Liu, W. (2022). Intellectual property protection for deep learning models: Taxonomy, methods, attacks, and evaluations. IEEE Transactions on Artificial Intelligence, 3(6), 908–923.
[27] Li, Y., Wang, H., & Barni, M. (2021). A survey of deep neural network watermarking techniques. Neurocomputing, 461, 171–193.
[28] Cao Y, Li S, Liu Y, et al. A survey of AI-generated content (AIGC). ACM Computing Surveys, 2025, 57(5): 1-38.
[29] Chen Y, Vice J, Akhtar N, Haldar N A, Mian A. Image watermarking of generative diffusion models. arXiv preprint arXiv:2502.10465, 2025.
[30] Chen Y, Akhtar N, Haldar N A, Mian A. Dynamic watermarks in images generated by diffusion models. arXiv preprint arXiv:2502.08927, 2025.
[31] Chen Y, Ma Z, Fang H, Zhang W, Yu N. TAG-WM: Tamper-aware generative image watermarking via diffusion inversion sensitivity. arXiv preprint arXiv:2506.23484, 2025.
[32] Jovanović N, Labiad I, Souček T, Vechev M, Fernandez P. Watermarking autoregressive image generation. arXiv preprint arXiv:2506.16349, 2025.
[33] Coalition for Content Provenance and Authenticity. C2PA technical specification: content credentials for digital provenance. C2PA, 2025.
[34] Zhang R, Li X, Wang Y, et al. A flexible and privacy-preserving cross-chain identity authentication system based on anonymous credentials. Proceedings of the ACM Conference on Data and Application Security and Privacy, 2024.
[35] Chen Z, Liu H, Zhang L, Dai B, Shi Y. Research on key technologies for privacy-preserving, regulatorily compliant, and cross-chain interoperability in

heterogeneous blockchain systems. Scientific Reports, 2026, 16: 12817.
[36] Yu H, Chen Y, Su S, Su J, Chen Y, Yang Z. DART: Distributed Zero Knowledge Data Auditing With Retrievability for Blockchain-Based Decentralized Storage Networks. IEEE Transactions on Information Forensics and Security, 2025, 20: 11264-11278.
[37] Jin Z, Zhang X, Su J, Zhang L, Shen J. Subgraph-Driven Lightweight Federated Learning for Spatiotemporal Cellular Traffic Prediction. IEEE Transactions on Network and Service Management, 2026, 23: 1435-1448.
[38] Jin Z, Yang C, Ye Y, Zhang L, Shen J, Su J. Mobility-Aware Semi-Asynchronous Federated Learning for Vehicular Networks. IEEE Transactions on Vehicular Technology, 2026, 75(2): 2001-2012.
[39] Su J, Jiang M. A Hybrid Entropy and Blockchain Approach for Network Security Defense in SDN-Based IIoT. Chinese Journal of Electronics, 2023, 32(3): 531-541.
[40] Xia C, Jin Z, Su J, Li B. Mobility-Aware Offloading and Resource Allocation Strategies in MEC Network Based on Game Theory. Wireless Communications and Mobile Computing, 2023, 2023: 5216943.