Journal of Information Systems Engineering and Management
2026, 11(1s)
e-ISSN: 2468-4376
https://www.jisem-journal.com

Engineering Scalable and Compliant Payment Systems for Autonomous Agentic Commerce

Arpit Mittal
IEEE Senior, USA

ARTICLE INFO

Received: 03 Dec 2025

Revised: 10 Jan 2026

Accepted: 20 Jan 2026

ABSTRACT

Autonomous agentic commerce represents a fundamental transformation in digital transaction ecosystems. Software agents executing commercial functions on behalf of human principals introduce unprecedented challenges for payment infrastructure. Traditional payment architectures designed for human-initiated transactions cannot accommodate machine-speed processing requirements. Existing compliance frameworks presuppose direct human accountability at transaction origination points. The absence of integrated architectural solutions addressing scalability, regulatory compliance, identity management, and dispute resolution simultaneously creates significant barriers to agentic commerce adoption. This article presents a comprehensive framework for payment system architecture supporting autonomous commercial agents. The Verifiable Agent Credential (VAC) protocol establishes cryptographic binding between agent identities and human principal profiles. Agent attestation services validate authorization scope against verified identity profiles for each transaction. Real-time sanctions screening employs phonetic matching algorithms and transliteration logic supporting multiple script systems. Behavioral fraud detection models calibrated for non-human transaction patterns identify compromised credentials and account takeover attempts. Automated dispute resolution mechanisms generate comprehensive evidentiary timelines enabling liability attribution across delegation relationships. The framework synthesizes distributed systems architecture, cryptographic protocols, and machine learning (ML) capabilities into a cohesive design paradigm. Payment infrastructure incorporating these architectural principles supports the projected expansion of agentic commerce while maintaining regulatory adherence and consumer protection guarantees essential to financial system integrity.

Keywords: Autonomous Agentic Commerce, Payment System Architecture, Verifiable Agent Credentials, Regulatory Compliance Framework, Behavioral Fraud Detection, Automated Dispute Resolution

I. Introduction

Autonomous agents are increasingly assuming commercial functions previously exclusive to human actors. These software entities execute purchasing decisions, optimize expenditure patterns, and orchestrate payment flows on behalf of consumers and merchants alike. Consumer-facing agents function as delegated shoppers, while merchant-integrated agents dynamically adjust pricing strategies and manage inventory replenishment.

Small businesses and independent merchants derive particular advantage from agentic payment infrastructure through substantially simplified integration pathways. Traditional payment system onboarding requires extensive technical expertise, prolonged configuration procedures, and significant capital investment. These barriers disproportionately affect smaller commercial entities. Autonomous integration agents abstract this complexity through single-interaction deployment mechanisms, enabling merchants to establish compliant payment acceptance capabilities without dedicated technical personnel. This streamlined onboarding democratizes access to sophisticated payment infrastructure previously available only to large enterprises with substantial information

technology resources. The reduction of integration timelines from weeks to minutes fundamentally alters competitive dynamics within retail and service sectors.

The mass adoption of agentic artificial intelligence (AI) raises significant ethical concerns regarding accountability, transparency, and decision-making processes [1]. Autonomous systems operating without continuous human oversight introduce questions of liability attribution. The delegation of commercial authority to non-human actors fundamentally alters established transactional relationships. Beyond accountability, transparency concerns emerge when agents execute transactions through opaque decision-making processes that principals cannot readily audit or understand. Ensuring that human principals maintain meaningful oversight of agent actions remains a critical design consideration for payment infrastructure.

Contemporary payment infrastructure reveals fundamental inadequacies when confronted with agentic commerce demands. Existing systems optimize for human cognitive and behavioral constraints. Transaction velocities measured in seconds reflect human decision-making timelines. Fraud detection patterns presuppose human behavioral characteristics. Dispute resolution mechanisms depend upon direct human involvement. Agentic commerce invalidates each of these foundational assumptions. Transaction volumes generated by autonomous agents exceed current capacity by substantial margins. Behavioral patterns exhibited by software agents differ markedly from human norms. Liability complexities emerge from authority structures lacking clear precedent in financial regulation.

Machine intelligence applications within commercial environments demonstrate considerable potential for transforming market operations. Advanced data processing capabilities enable autonomous systems to analyze transaction patterns and predict market trends with notable accuracy [2]. Large-scale data processing architectures support the computational demands of continuous commercial monitoring. Trend prediction mechanisms inform agent decision-making across purchasing, pricing, and inventory management functions. The integration of ML within commercial workflows accelerates transaction throughput beyond human-manageable velocities. This acceleration necessitates corresponding evolution in payment infrastructure design.

The ethical dimensions of agentic commerce extend beyond technical performance considerations to encompass fairness, bias mitigation, and equitable treatment. Questions of accountability arise when autonomous agents execute transactions resulting in financial harm [1]. Attribution of responsibility between human principals and software agents lacks established legal frameworks. Consumer protection mechanisms designed for human-to-human commerce require adaptation. The proliferation of agentic commerce introduces novel fraud vectors exploiting the gap between human intent and agent interpretation. Payment systems must therefore incorporate mechanisms for establishing clear chains of accountability while ensuring algorithmic fairness across diverse user populations.

The research gap addressed herein concerns the absence of integrated architectural frameworks. Current approaches fail to simultaneously achieve hyper-scalability for continuous micro-transactions, real-time compliance verification across jurisdictional boundaries, cryptographically verifiable identity binding between agents and principals, and automated liability attribution for dispute resolution. Existing literature addresses these challenges in isolation. Distributed systems research optimizes throughput without regulatory consideration. Compliance frameworks presuppose human-speed transaction flows. Identity management protocols lack provisions for agent authority. This article contributes a comprehensive design paradigm addressing these interdependent requirements through novel protocol specifications and system architectures.

Journal of Information Systems Engineering and Management
2026, 11(1s)
e-ISSN: 2468-4376
https://www.jisem-journal.com

II. Related Work and Methodology

Existing literature addresses payment system scalability and regulatory compliance as isolated concerns. Distributed systems research optimizes transaction throughput without consideration for compliance constraints. Regulatory technology frameworks assume human-speed transaction flows incompatible with agentic commerce velocities. Identity management protocols lack provisions for authority relationships between autonomous agents and human principals. Fraud detection models trained exclusively on human behavioral patterns fail to identify agent-specific anomaly signatures. This fragmentation prevents coherent architectural design for production-scale agentic commerce deployments.

The methodology underlying the proposed framework synthesizes principles from multiple technical domains. Distributed event-streaming architectures provide foundational patterns for scalable transaction processing. Cryptographic attestation mechanisms inform the VAC protocol design. Phonetic matching algorithms and transliteration techniques drawn from information retrieval support multilingual sanctions screening. ML classification approaches enable behavioral fraud detection calibrated for non-human transaction patterns. Automated evidence collection techniques from digital forensics support dispute resolution timeline generation.

The framework contributes an integrated architectural paradigm addressing interdependent requirements simultaneously. Novel protocol specifications establish cryptographic binding between agent and principal identities. System architectures enable independent scaling of processing, compliance, and fraud detection components. The synthesis of established techniques within a unified design enables payment infrastructure supporting agentic commerce at scale.

III. Scalable Transaction Processing Architecture

A. Throughput Optimization for Non-Human Actors

Traditional payment processing architectures impose latency characteristics acceptable for human-initiated transactions. This latency proves prohibitive for autonomous agent operations. Agentic commerce demands near-instantaneous authorization responses. Continuous high-frequency transaction streams require processing without cascading failures. Service quality degradation undermines the viability of agentic commerce. The architectural challenge extends beyond simple throughput increases and requires fundamental redesign of processing pipelines.

Modern e-commerce platform development emphasizes modular architectures capable of handling variable transaction loads. Web-based platforms serving telecommunications and retail sectors demonstrate the importance of scalable backend infrastructure [3]. Database optimization and caching mechanisms reduce response latencies for high-frequency requests. Load balancing distributes transaction processing across multiple server instances. These architectural patterns translate directly to payment system requirements for agentic commerce. The separation of presentation layers from business logic enables independent scaling of transaction processing components.

Distributed event-streaming mechanisms provide architectural foundations for scalable payment processing. Event-driven designs enable asynchronous processing of transaction requests. Horizontal scaling accommodates variable workload demands without service interruption. Authorization logic operates independently from compliance verification modules. Fraud assessment functions execute in parallel with core transaction processing. This decomposition prevents bottlenecks within any single processing domain. Payment systems benefit substantially from this architectural flexibility. Peak transaction periods trigger automatic resource allocation, while reduced activity periods allow resource consolidation.

Journal of Information Systems Engineering and Management
2026, 11(1s)
e-ISSN: 2468-4376
https://www.jisem-journal.com

B. Dynamic Policy Evaluation

Compliance requirements vary substantially across jurisdictions. Regulatory frameworks evolve constantly in response to emerging risks. Static compliance implementations become outdated rapidly. System redeployment for policy updates introduces operational risk. The architecture therefore incorporates policy engines capable of dynamic rule evaluation. Transaction assessment occurs against continuously updated rule sets. Regulatory changes propagate through the system without requiring code modifications.

AI and data analytics capabilities transform compliance monitoring within payment systems. ML algorithms identify patterns indicative of regulatory violations across large transaction datasets [4]. Real-time analytics enable immediate detection of suspicious transaction sequences. Competitive intelligence frameworks demonstrate the value of continuous data processing for identifying emerging risks. These analytical approaches apply directly to compliance monitoring requirements. Pattern recognition across transaction streams identifies potential violations before completion.

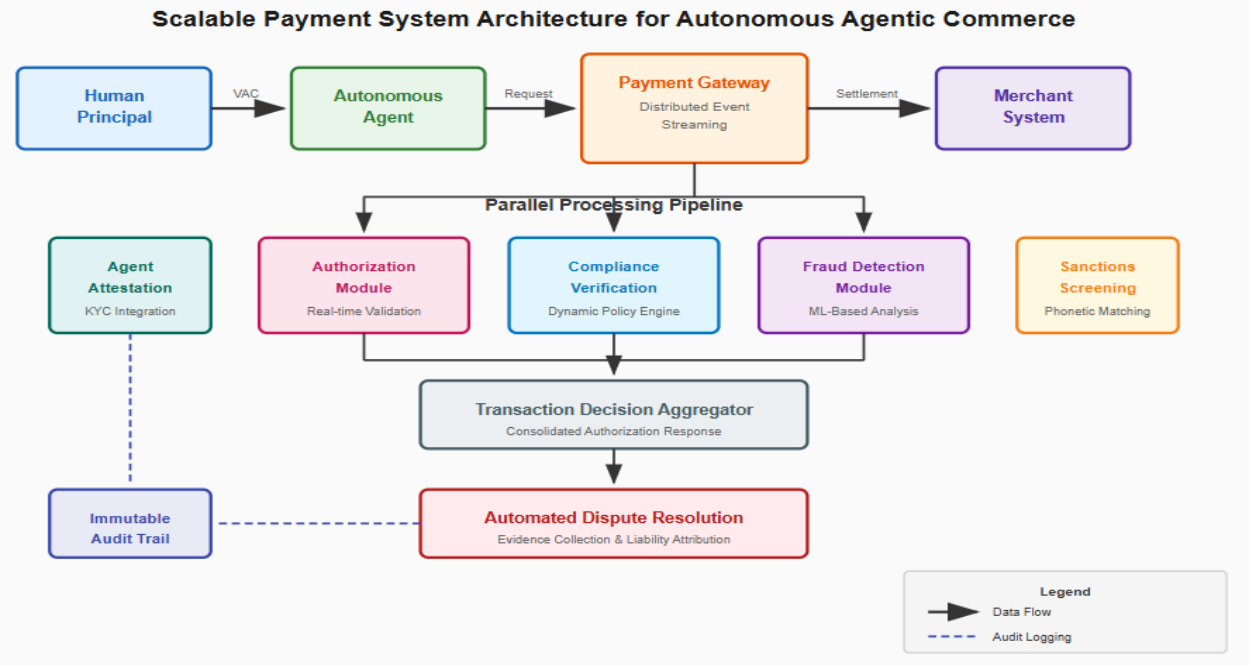

Fig 1. System Architecture Overview [3, 4].

Policy definitions utilize declarative specifications for accessibility and maintainability. Non-technical compliance personnel modify evaluation criteria through structured interfaces. System integrity remains protected through validation mechanisms. Audit controls document all policy modifications for regulatory review. The separation of policy logic from processing infrastructure enables rapid regulatory response. New compliance requirements enter production within abbreviated timeframes. Dynamic policy evaluation supports the jurisdictional complexity inherent in cross-border agentic commerce.

Component	Function	Characteristic
Distributed Event-Streaming	Asynchronous transaction processing	Horizontal scalability
Parallel Pipeline Stages	Separation of processing functions	Independent resource scaling
Authorization Logic Module	Transaction validation	Real-time response capability

Compliance Verification Module	Regulatory rule evaluation	Dynamic policy updates
Fraud Assessment Module	Anomaly detection	Parallel execution
Load Balancing	Workload distribution	Multiple server instances
Policy Engine	Declarative rule processing	Non-technical modification support

Table 1. Distributed Processing Mechanisms for Agentic Commerce [3, 4].

III. Regulatory Compliance Framework

A. Agent Attestation Services

Each transaction initiated by an autonomous agent undergoes programmatic validation. The validation confirms authorization scope against the human principal's verified identity profile. This attestation mechanism ensures agentic commerce activity remains within established boundaries. Identity verification processes define these boundaries during initial principal onboarding. Immutable audit trails document authorization state at transaction time. The attestation framework addresses fundamental questions of accountability in agentic commerce.

Digital identity verification systems provide foundational infrastructure for agent attestation services. Know-Your-Customer (KYC) protocols establish verified identity profiles for human principals. Document verification and biometric authentication confirm principal identities during onboarding. Agent credentials derive from these verified principal profiles through cryptographic binding mechanisms. The delegation relationship between principal and agent requires continuous verification. Authorization scope limitations prevent agents from exceeding granted commercial authorities. Each attestation event generates cryptographic proof of authorization validity.

Audit trail generation represents a critical component of attestation services. Each transaction produces immutable records documenting the authorization context. Principal identity verification status receives permanent documentation. Agent credential validity and scope limitations appear within audit records. Regulatory examination requires access to comprehensive transaction histories. Attestation services maintain detailed logs supporting compliance demonstration. The evidentiary value of these records depends upon integrity and completeness.

B. Real-Time Sanctions Screening

Compliance with sanctions regulations requires screening transaction participants against restricted party lists. Regulatory authorities maintain these lists across multiple jurisdictions. Transaction processing cannot proceed until sanctions verification completes. Real-time screening demands optimized matching algorithms and efficient data structures. The screening architecture must accommodate frequent list updates without service interruption.

Distributed detection architectures offer valuable design patterns for sanctions screening systems. Network intrusion detection systems demonstrate effective approaches for real-time pattern matching across high-volume data streams [5]. Signature-based detection methods identify known threat patterns through optimized matching algorithms. These detection architectures distribute processing loads across multiple nodes for scalability. The application of distributed detection principles to sanctions screening enables real-time verification at agentic transaction velocities. Pattern matching against restricted party signatures occurs in parallel across processing clusters.

Phonetic matching algorithms address name variation challenges in sanctions screening. Double Metaphone algorithms encode names according to pronunciation patterns. Soundex variants provide

alternative encoding schemes for different linguistic contexts. Transliteration logic extends matching capabilities to non-Latin scripts. Arabic name variations require specialized handling due to transliteration inconsistencies. Cyrillic and Thai character systems present similar challenges. High-accuracy matching depends upon comprehensive linguistic coverage.

Text mining and automated analysis techniques enhance sanctions screening effectiveness. Research on automated text mining demonstrates sophisticated approaches for pattern identification across large datasets [6]. These techniques extract meaningful patterns from unstructured textual data. Name matching benefits from similar analytical approaches. The evolution of algorithmic methods provides increasingly accurate matching capabilities. Computational efficiency improvements enable real-time processing of complex matching operations. The balance between matching accuracy and processing speed defines screening system effectiveness within agentic commerce environments.

Mechanism	Purpose	Implementation
Agent Attestation Service	Authorization scope validation	Cryptographic binding verification
KYC Profile Integration	Principal identity verification	Document and biometric authentication
Immutable Audit Trails	Transaction documentation	Timestamped authorization records
Double Metaphone Algorithm	Phonetic name matching	Pronunciation pattern encoding
Soundex Variants	Alternative name encoding	Linguistic context adaptation
Transliteration Logic	Non-Latin script support	Arabic, Cyrillic, and Thai processing
Distributed Detection Architecture	Pattern matching scalability	Parallel processing clusters

Table 2. Agent Attestation and Sanctions Screening Components [5, 6].

IV. Identity Binding and Fraud Mitigation

A. Verifiable Agent Credential Protocol

The VAC protocol establishes cryptographic binding between autonomous agent identities and human principal profiles. This mechanism employs tokenization techniques to ensure credential security. Agent credentials cannot be transferred between principals. Duplication attempts trigger immediate credential revocation. Exploitation independent of the authorizing principal becomes technically infeasible. The protocol architecture addresses fundamental authentication challenges in agentic commerce.

Traditional delegation mechanisms present significant security limitations for agentic commerce. Open Authorization (OAuth) frameworks assume human involvement in authentication flows. Token refresh procedures depend upon interactive principal participation. Autonomous agents operating continuously cannot rely upon these interactive mechanisms. The VAC protocol eliminates interactive dependencies. Cryptographic binding ensures continuous authentication without human intervention. Non-repudiation guarantees exceed those achievable through traditional approaches.

Attack vector mitigation represents a primary design consideration within the protocol. Agent impersonation attempts require possession of cryptographic keys. Mandate exploitation through

Journal of Information Systems Engineering and Management
2026, 11(1s)
e-ISSN: 2468-4376
https://www.jisem-journal.com

scope expansion triggers validation failures. Formal security properties define protocol resistance against known attack categories. Credential revocation propagates immediately across all verification points. Compromised credentials become unusable before exploitation occurs. The binding between agent identity and principal profile remains continuously verifiable.

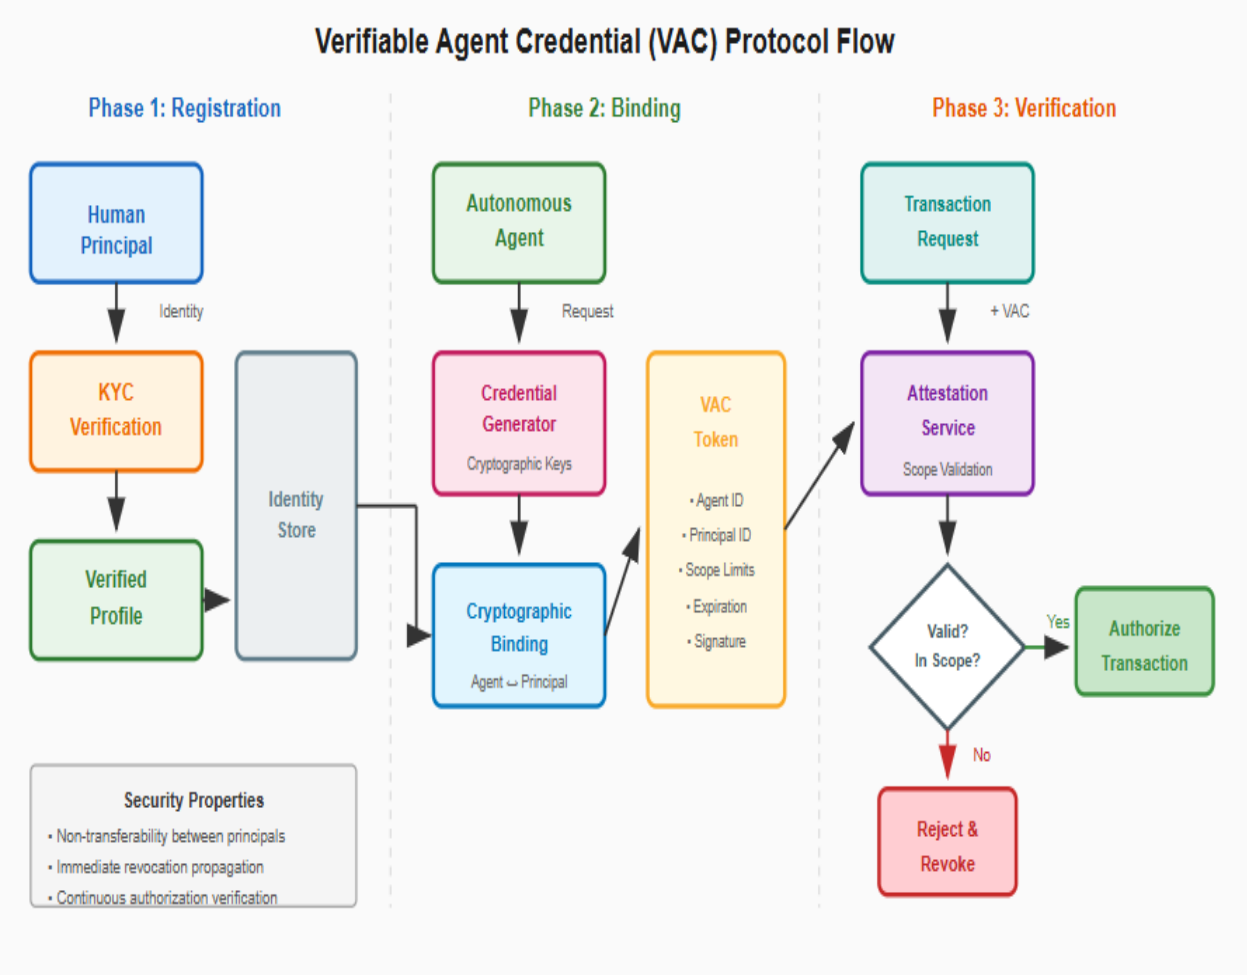

Fig 2. Verifiable Agent Credential Protocol

B. Behavioral Fraud Detection

Fraud detection models trained on human behavioral patterns exhibit diminished effectiveness for agent-initiated transactions. Human transaction patterns reflect cognitive limitations and decision-making characteristics. Autonomous agents generate fundamentally different transaction patterns. Structured Application Programming Interface (API) interaction sequences replace variable human input patterns. Consistent transaction velocities differ from human timing variations. The framework incorporates ML models calibrated for non-human behavioral characteristics.

Comparative analysis of ML models reveals varying effectiveness for fraud detection applications. Research on credit card fraud detection demonstrates that ensemble methods and gradient boosting algorithms provide superior classification accuracy [7]. Random forest classifiers effectively identify fraudulent patterns within imbalanced transaction datasets. Support vector machines (SVMs) offer robust performance for high-dimensional feature spaces. The selection of appropriate algorithms depends upon specific fraud pattern characteristics. Agent-initiated fraud presents distinct pattern signatures requiring specialized model configurations.

Journal of Information Systems Engineering and Management
2026, 11(1s)
e-ISSN: 2468-4376
https://www.jisem-journal.com

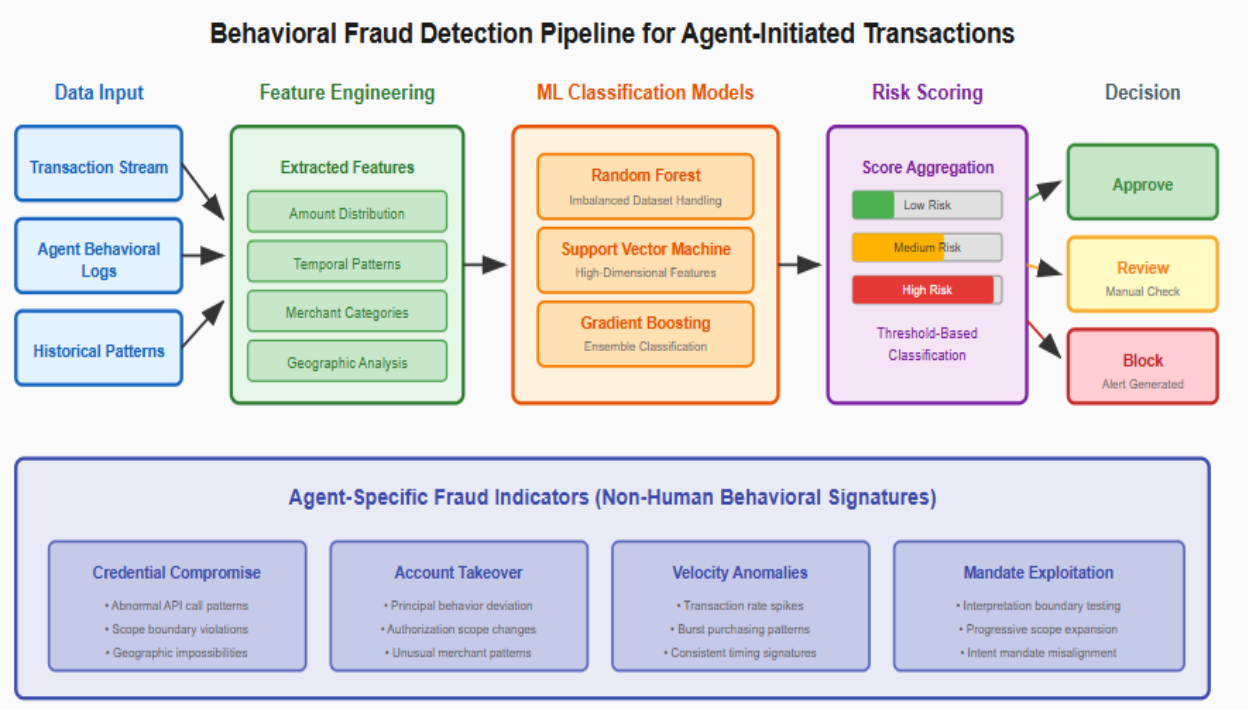

Fig 3. Behavioral Fraud Detection Pipeline [7, 8].

Feature engineering plays a critical role in fraud detection model performance. Transaction amount distributions provide baseline behavioral indicators. Time-based features capture temporal patterns within transaction sequences. Merchant category patterns reveal purchasing behavior anomalies [8]. Geographic location analysis identifies physically impossible transaction sequences. The combination of multiple feature categories enhances detection accuracy. Agent-specific features extend traditional approaches to accommodate non-human behavioral characteristics.

Model Performance Metrics (n=1,247,583 transactions, 2022-2024)

CLASSIFICATION RESULTS

	Predicted: Fraudulent	Predicted: Legitimate	
Actual **Fraudulent**	4,847 (True Pos)	173 (False Neg)	Sensitivity: 96.55%
Legitimate	2,614 (False Pos)	1,239,949 (True Neg)	Specificity: 99.79%

Overall Accuracy: 98.76%
Precision: 64.96%
F1 Score: 77.64%
AUC-ROC: 0.9923

Top Predictive Features (Importance Score):

1.	Transaction velocity deviation index	0.194
2.	Agent credential scope boundary proximity	0.167
3.	Geographic impossibility score	0.143
4.	Merchant category deviation ratio	0.128
5.	Principal behavioral baseline divergence	0.112
6.	API call pattern anomaly coefficient	0.097
7.	Transaction amount distribution skewness	0.084
8.	Temporal pattern consistency index	0.075

Fig 4. Predictive Model Performance for Agent-Initiated Fraud Detection Classification (Illustrative). [7, 8].

C. Ethical Considerations in Fraud Detection

The deployment of ML-based fraud detection within agentic commerce raises significant ethical considerations requiring explicit attention. Algorithmic bias represents a primary concern, as models trained on historical transaction data may perpetuate or amplify existing discriminatory patterns. Fraud detection systems that disproportionately flag transactions from specific demographic groups, geographic regions, or merchant categories introduce fairness concerns requiring systematic mitigation.

Transparency in fraud detection decision-making enables meaningful human oversight of automated systems. Black-box models that flag transactions without interpretable justification undermine principal trust and complicate dispute resolution. The framework therefore incorporates explainability mechanisms that generate human-readable rationales for fraud determinations. These rationales support appeal processes and enable continuous model improvement through feedback integration.

Bias mitigation strategies include regular algorithmic audits examining detection rates across demographic categories, geographic distributions, and transaction types. Training data curation processes must identify and address historical biases that could propagate into deployed models. Fairness constraints incorporated during model training ensure equitable treatment across protected categories. Continuous monitoring of deployed models detects emergent bias patterns requiring corrective intervention.

Model training requires careful attention to dataset characteristics. Class imbalance between legitimate and fraudulent transactions presents significant challenges. Oversampling techniques address minority class representation within training data. Cross-validation procedures ensure model generalization across unseen transactions [7]. Continuous model retraining incorporates newly identified fraud patterns. The dynamic nature of fraud techniques necessitates adaptive detection capabilities. Agent credential compromise and account takeover attempts generate identifiable behavioral signatures within properly calibrated detection systems.

Element	Description	Security Property
Verifiable Agent Credential	Cryptographic identity binding	Non-transferability
Tokenization Techniques	Credential protection	Duplication prevention
Credential Revocation	Compromise response	Immediate propagation
Random Forest Classifier	Fraud pattern identification	Imbalanced dataset handling
Support Vector Machine	High-dimensional classification	Robust performance
Ensemble Methods	Combined model approach	Superior classification accuracy
Feature Engineering	Behavioral indicator extraction	Transaction pattern analysis

Table 3. Verifiable Agent Credential Protocol and Detection Models [7, 8].

V. Automated Dispute Resolution Mechanisms

Dispute resolution within agentic commerce requires mechanisms capable of establishing liability attribution. Delegation relationships introduce complexity absent from traditional commerce disputes. Human principals express intent mandates through various interface modalities. Autonomous agents interpret these mandates according to programmed logic. Transaction executions

Journal of Information Systems Engineering and Management
2026, 11(1s)
e-ISSN: 2468-4376
https://www.jisem-journal.com

reflect agent interpretations rather than direct human actions. The framework implements automated evidence collection for comprehensive dispute analysis.

New technologies fundamentally transform dispute resolution processes across commercial domains. AI and blockchain technologies introduce unprecedented capabilities for arbitration and conflict management [9]. Smart contracts enable automated execution of resolution outcomes based upon predefined conditions. ML algorithms classify disputes according to established taxonomic categories. Natural language processing (NLP) extracts relevant factual assertions from dispute submissions. This technological integration represents a significant departure from traditional manual resolution procedures. The application of these advances to agentic commerce disputes enables scalable resolution capabilities.

Timeline generation documents the complete transaction lifecycle for evidentiary purposes. Intent mandates expressed by human principals receive timestamped recording. Agent interpretation processes produce logged decision traces. Transaction execution details appear within immutable audit records. The evidentiary foundation enables systematic liability assignment across multiple party categories. Principal instruction failures differ fundamentally from agent interpretation errors. Merchant fulfillment deficiencies represent a third distinct liability category. Appropriate chargeback allocation depends upon accurate categorization of fault origin.

Digital banking transformation provides contextual frameworks for automated dispute handling. Industry 4.0 principles emphasize automation and intelligent decision-making across financial services [10]. Digital banking implementations incorporate advanced technologies for customer service optimization. Dispute resolution represents a critical customer service function requiring technological enhancement. Multi-criteria decision-making approaches support complex liability determinations. The evaluation of multiple factors simultaneously enables nuanced attribution reflecting actual circumstances.

Evidence collection automation eliminates manual investigation requirements for routine disputes. Transaction logs provide objective records of agent actions throughout commercial interactions. Communication records document principal-agent exchanges preceding disputed transactions. Merchant response data confirms fulfillment status and delivery outcomes. Integration across disparate data sources enables comprehensive dispute reconstruction. Automated systems analyze collected evidence against established liability criteria without human intervention.

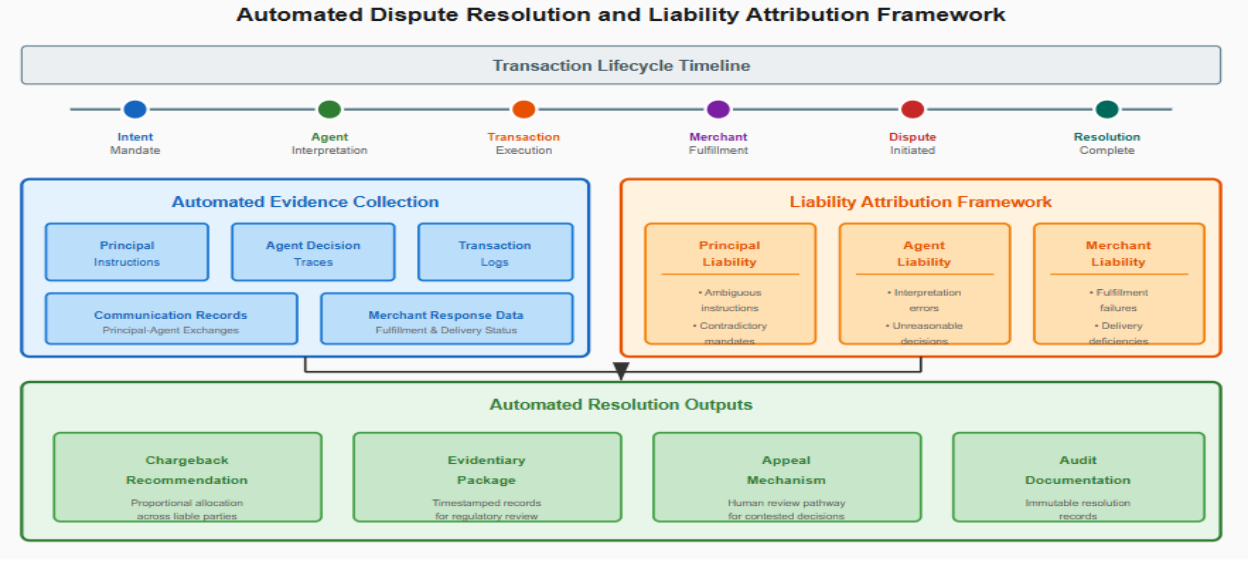

Fig 5. Automated Dispute Resolution Mechanism [9, 10].

A. Fairness in Automated Resolution

Automated dispute resolution mechanisms must incorporate fairness guarantees ensuring equitable treatment across all participants. Algorithmic determination of liability raises concerns regarding systematic bias favoring particular party categories. Merchants with extensive transaction histories may receive preferential treatment compared to occasional sellers. Principals with higher account values may experience different resolution outcomes than lower-value participants. The framework addresses these concerns through explicit fairness constraints within resolution algorithms.

Procedural fairness requires that all parties receive equivalent opportunities to present evidence and contest automated determinations. The framework implements standardized evidence submission interfaces ensuring consistent documentation requirements across party categories. Appeal mechanisms provide recourse for participants dissatisfied with automated outcomes. Human review remains available for complex or contested cases requiring judgment beyond algorithmic capability.

Outcome fairness demands that similar disputes receive similar resolutions regardless of participant characteristics unrelated to fault determination. Regular audits examine resolution patterns across merchant categories, transaction values, and geographic distributions. Statistical analysis identifies systematic deviations from expected outcome distributions requiring corrective intervention. Transparency reports document aggregate resolution statistics enabling external accountability assessment.

The liability attribution framework addresses unique challenges inherent in agentic commerce. Principal responsibility applies when intent mandates contain ambiguous or contradictory instructions. Agent liability emerges from interpretation errors inconsistent with reasonable understanding. Merchant accountability covers fulfillment failures independent of ordering accuracy. Hybrid liability scenarios require proportional allocation across multiple parties. The framework supports nuanced attribution reflecting actual fault distribution across delegation chains.

Chargeback processing within agentic commerce requires specialized handling procedures distinct from traditional approaches. Agent-initiated transactions introduce intermediary accountability questions absent from human-initiated commerce. Payment network rules require adaptation for agentic commerce scenarios. The automated resolution framework generates chargeback recommendations based on comprehensive evidence analysis. Appeal mechanisms provide recourse for disputed automated determinations.

Component	Function	Output
Timeline Generation	Transaction lifecycle documentation	Timestamped event records
Intent Mandate Recording	Principal instruction capture	Logged decision traces
Agent Interpretation Logging	Decision process documentation	Audit trail entries
Evidence Correlation	Multi-source data integration	Comprehensive reconstruction
Liability Classification	Fault category assignment	Attribution determination
Principal Liability	Ambiguous instruction identification	Proportional allocation
Agent Liability	Interpretation error detection	Fault distribution
Merchant Liability	Fulfillment failure assessment	Chargeback recommendation

Table 4. Liability Attribution and Evidence Collection Mechanisms [9, 10].

Journal of Information Systems Engineering and Management
2026, 11(1s)
e-ISSN: 2468-4376
https://www.jisem-journal.com

VII. Limitations and Future Considerations

A. Cryptographic Trust Dependencies

The VAC protocol relies upon cryptographic mechanisms that introduce specific trust dependencies. Public key infrastructure (PKI) underlying credential verification requires trusted certificate authorities. Compromise of root certificates could undermine system-wide credential validity. Key management procedures impose operational burdens upon both principals and system operators. Cryptographic algorithm obsolescence necessitates periodic protocol updates as computational capabilities advance. Quantum computing developments may require migration to post-quantum cryptographic primitives within foreseeable timeframes.

B. Computational Overhead

Real-time compliance verification, sanctions screening, and fraud detection impose substantial computational requirements. Phonetic matching algorithms processing multilingual names against extensive restricted party lists demand significant processing resources. ML-based fraud detection models require continuous inference operations across high-volume transaction streams. The aggregate computational overhead may introduce latency inconsistent with agentic commerce velocity requirements under peak load conditions. Resource scaling to accommodate worst-case scenarios incurs significant infrastructure costs.

C. Adoption Barriers

Widespread agentic commerce adoption faces multiple barriers beyond technical architecture. Regulatory frameworks governing autonomous commercial agents remain underdeveloped across most jurisdictions. Legal liability attribution between principals and agents lacks established precedent. Consumer acceptance of agent-mediated commerce depends upon trust development through demonstrated reliability. Merchant integration requires adaptation of existing systems to accommodate agent-initiated transactions. Payment network rule modifications necessary for agentic commerce processing require industry-wide coordination.

D. Interoperability Challenges

The proposed framework assumes deployment within controlled infrastructure environments. Integration with legacy payment systems operating on incompatible protocols presents significant challenges. Cross-network agent credential verification requires standardization efforts extending beyond individual implementations. International agentic commerce spanning multiple payment networks and regulatory jurisdictions introduces coordination complexities not fully addressed within the current framework.

E. Adversarial Adaptation

Fraud detection models calibrated for current agent behavioral patterns face adversarial adaptation challenges. Malicious actors observing detection mechanisms may modify attack strategies to evade identification. The continuous retraining requirements introduce delays during which novel attack patterns remain undetected. Sophisticated adversaries may exploit the gap between fraud pattern emergence and model adaptation. Maintaining detection effectiveness requires ongoing investment in threat intelligence and model development.

Conclusion

The emergence of autonomous commercial agents necessitates fundamental reconceptualization of payment infrastructure design principles. Software entities executing commercial functions operate at velocities incompatible with legacy system architectures. Regulatory frameworks developed for

human-mediated commerce require adaptation for machine-initiated transactions. The architectural framework presented within this article addresses interconnected challenges spanning transaction processing, compliance verification, identity management, fraud prevention, and dispute resolution domains.

Distributed event-streaming mechanisms enable horizontal scaling for high-frequency transaction streams. Dynamic policy evaluation engines accommodate evolving regulatory requirements across jurisdictional boundaries without service interruption. The VAC protocol provides cryptographic guarantees exceeding traditional delegation mechanisms. Agent attestation services maintain continuous authorization verification throughout transaction lifecycles. Sanctions screening architectures employ sophisticated matching algorithms supporting diverse linguistic systems. ML models trained on non-human behavioral patterns detect agent-specific fraud indicators while incorporating bias mitigation strategies ensuring equitable treatment. Automated evidence collection generates comprehensive transaction timelines supporting liability attribution without manual investigation, with fairness constraints ensuring consistent resolution outcomes.

The integration of these capabilities establishes foundational infrastructure for emerging agentic commerce ecosystems. Payment systems incorporating the architectural principles outlined enable scalable agentic commerce while preserving regulatory compliance and consumer protection requirements. Future development directions include formal protocol verification, empirical evaluation across deployment contexts, adaptation for emerging regulatory frameworks governing autonomous commercial entities, and continued refinement of ethical safeguards ensuring transparency, fairness, and accountability throughout agentic commerce operations.

References

[1] Sonja Jovanović et al., "ETHICAL CONCERNS AND MASS AGENTIC AI ADOPTION," SINTEZA 2025. [Online]. Available: https://portal.sinteza.singidunum.ac.rs/Media/files/2025/369-375.pdf

[2] Olalekan Hamed Olayinka, "Revolutionizing market analysis using machine intelligence, trend prediction, and large-scale data processing," World Journal of Advanced Research and Reviews, 2023. [Online]. Available: https://www.researchgate.net/profile/Olayinka-Olalekan/publication/390150074

[3] Adrian-Costin MINEA et al., "Development of a Web-Based e-Commerce Platform for a Telecommunications Provider," Sciendo, 2025. [Online]. Available: https://reference-global.com/2/v2/download/pdf/10.2478/picbe-2025-0380

[4] Harper Jack, "Leveraging AI and Data Analytics: Revolutionizing Competitive Intelligence for Market Insights," ResearchGate, 2024. [Online]. Available: https://www.researchgate.net/profile/Harper-Jack/publication/387130618_Leveraging_AI_and_Data_Analytics_Revolutionizing_Competitive_Intelligence_for_Market_Insights/links/6761840da3978e15e7905d02/Leveraging-AI-and-Data-Analytics-Revolutionizing-Competitive-Intelligence-for-Market-Insights.pdf

[5] S.Subapriya and Ms.N.Radhika, "DNIDPS: Distributed Network Intrusion Detection and Prevention System," International Journal of Innovative Science, Engineering & Technology, 2014. [Online]. Available: https://www.ijiset.com/v1s7/IJISET_V1_I7_11.pdf

[6] Anil Jadhav et al., "Evolution of Software Development Effort and Cost Estimation Techniques: Five Decades Study Using Automated Text Mining Approach," Mathematical Problems in Engineering, 2022. [Online]. Available: https://onlinelibrary.wiley.com/doi/pdf/10.1155/2022/5782587

[7] Masad A. Alrasheedi, "Enhancing Fraud Detection in Credit Card Transactions: A Comparative Study of Machine Learning Models," Computational Economics, 2025. [Online]. Available: https://link.springer.com/content/pdf/10.1007/s10614-025-11071-3.pdf

[8] Akshat Shah and Yogeshvari Makwana, "Credit Card Fraud Detection," ResearchGate. [Online]. Available: https://www.researchgate.net/profile/Akshat-Shah-27/publication/369857378_Credit_Card_Fraud_Detection/links/642fbd0d20f25554da158af5/Credit-Card-Fraud-Detection.pdf

[9] Magdalena Łągiewska, "New Technologies in International Arbitration: A Game-Changer in Dispute Resolution?" Springer, 2024. [Online]. Available: https://link.springer.com/content/pdf/10.1007/s11196-023-10070-7.pdf

[10] Maghsoud Amiri et al., "Evaluation of Digital Banking Implementation Indicators and Models in the Context of Industry 4.0: A Fuzzy Group MCDM Approach," MDPI, 2023. [Online]. Available: https://www.mdpi.com/2075-1680/12/6/516