Review

A Framework for Secure Communication in Decentralized AI Agent Systems

Mohammed Shakawat Hossen , Md Naeem Uddin , Md Zeaul Kader Opel , Shuvo Barua , Minhaz Uddin Piplu , Joynal Abedin Shahin , Md. Badiuzzaman Biplob *

Posted Date: 16 July 2025

doi: 10.20944/preprints202507.1162.v1

Keywords: decentralized AI agents; secure communication; cryptography; decentralized identity; multi-agent systems; threat detection; blockchain integration

 doi:10.20944/preprints202507.1162.v1

Review

A Framework for Secure Communication in Decentralized AI Agent Systems

Mohammed Shakawat Hossen [1], Md Naeem Uddin [2], Md Zeaul Kader Opel [3], Shuvo Barua [4], Minhaz Uddin Piplu [5], Joynal Abedin Shahin [6], and Md. Badiuzzaman Biplob [7,*]

Department of Computer Science and Engineering, International Islamic University Chittagong, Chattogram-4318, Bangladesh
* Correspondence: biplob.cse@iiuc.ac.bd

Abstract: This paper proposes a novel, layered security framework for decentralized AI agent communication. The framework integrates decentralized identity verification using Decentralized Identifiers (DIDs), hybrid encryption with optional post-quantum resilience, AI-inspired dynamic threat detection, and optional blockchain-based audit logging. Unlike prior solutions that address these components in isolation, this architecture unifies them into a cohesive, modular design that can be tailored to domain-specific requirements. A formal protocol specification is provided to support future formal verification and implementation. Although no empirical testing was conducted in this study, the framework offers a conceptual foundation for building trustworthy, adaptable communication systems for decentralized AI agents.

Keywords: Decentralized AI Agents, Secure Communication, Cryptography, Decentralized Identity, Multi-Agent Systems, Threat Detection, Blockchain Integration

---

1. Introduction

Decentralized AI agents are autonomous software entities that collaborate within distributed networks without reliance on a central coordinator. Such systems, often implemented as multi-agent systems (MAS), support scalability, resilience, and flexibility in domains like smart grids, autonomous logistics, healthcare IoT, and financial technology. However, the absence of central oversight introduces new security challenges. Agents must securely authenticate one another, protect data confidentiality and integrity during communication, detect malicious behavior, and support auditability — all in dynamic, often heterogeneous environments.

Ensuring secure communication among decentralized AI agents is especially challenging due to risks such as identity spoofing, man-in-the-middle attacks, data injection, and replay attacks. Traditional solutions like TLS/SSL depend on centralized certificate authorities and are ill-suited for fully decentralized architectures. Similarly, while blockchain-based solutions provide auditability, they can introduce latency and resource overhead that are impractical in time-sensitive or resource-constrained systems.

To address these gaps, this paper proposes a novel layered security framework that integrates decentralized identity verification using Decentralized Identifiers (DIDs), hybrid encryption (AES combined with optional post-quantum cryptography), AI-inspired anomaly detection, and optional blockchain-based logging for auditability. The architecture is modular and adaptable, enabling domain-specific customization. A formal protocol specification is also provided to guide future verification and implementation.

The rest of this paper is organized as follows: Section 2 reviews related work, Section 3 describes the proposed framework, Section 4 presents conceptual analysis and discussion, and Section 5 concludes with limitations and future work.

 doi:10.20944/preprints202507.1162.v1

2. Literature Review

A study by Huang et al. (2025) [1] introduces a novel zero-trust identity framework for agentic AI based on decentralized identifiers (DIDs), verifiable credentials, and policy-based access control. The framework emphasizes dynamic capability verification and anomaly detection at the identity layer. A strength is its alignment with decentralized governance models; however, broader interoperability with existing agent communication standards remains unexplored.

Ehtesham et al. (2025) [2] present a survey of emerging multi-agent communication standards—MCP, ACP, A2A, ANP—highlighting their strengths in enabling interoperability, secure discovery, and peer-to-peer messaging paradigms. While the protocols offer modular design patterns, the study calls for unified threat models across heterogeneous agent ecosystems.

Ranjan et al. (2025) [3] propose LOKA, a decentralized ethical agent protocol incorporating transparent governance modules alongside security mechanisms such as cryptographic binding and DID-based authentication. The work's novelty lies in ethical-integrity integration, although latency trade-offs in consensus for large-scale deployments are not fully evaluated.

A survey by Wang et al. (2025) [4] explores how blockchain frameworks support secure AI-agent collaboration, concluding that while immutability and smart contracts enhance trust, challenges remain around latency, throughput, and privacy in large-scale multi-agent systems.

Nagothu et al. (2022) [5] compare lightweight consensus algorithms for IoT agent networks, demonstrating that DAG-based protocols outperform simplistic blockchain in reducing communication overhead. The caveat is that such mechanisms may lack formal auditability compared to traditional blockchains.

Liu et al. (2023) [6] introduce an LSTM-based anomaly detector that identifies abnormal agent behaviors in swarm systems. The model adapts to sequential patterns but requires extensive runtime adaptation as agents learn new behaviors over time.

Chalvatzakis etal. (2024) [7] apply Isolation Forests to detect outlier communication events in industrial agent networks, showing that unsupervised anomaly detection can flag compromised nodes. However, the study notes false positives from benign irregularities and highlights the need for model recalibration.

Sharma et al. (2024) [8] combine DIDs with blockchain logging to manage IIoT security. Their hybrid protocol allows verifiable identity checks and decentralized audit trails, but the added complexity raises questions on scalability in large deployments.

A cooperative anomaly detection scheme presented by IJCAI 2023 participants [9] uses RNN predictors embedded in each agent to detect adversarial behavior within decentralized MAS. The decentralized architecture avoids reliance on central observers and improves resilience to targeted attacks.

Andreu (2025) [10] surveys decentralized agentic AI overviews, emphasizing the role of reputation systems and robust consensus in building secure multi-agent frameworks. The article frames security as intrinsic to agent architecture, not just cryptographic overlays.

Li et al. (2025) [11] propose SAFEFLOW, a protocol-level information-flow control framework for LLM-empowered agents. SAFEFLOW enforces confidentiality labels and transaction-level consistency, representing a leap beyond message encryption. Its practicality in dynamic MAS remains to be seen.

Hernandez-Ramos et al. (2023) [12] present an agent-centric security architecture that leverages DIDs at the agent-scope level. The design supports secure agent discovery and attribute-based control flows, offering a finer-grained model than entity-level identity, but the authors note a lack of evaluation in high-latency environments.

A study conducted by Pavle et al. (2025) [13] presents a federated Isolation Forest architecture for edge-based IoT systems, aimed at enabling decentralized anomaly detection among agents without sharing raw data. The authors demonstrate the system's potential to reduce communication overhead while maintaining detection fidelity, but acknowledge challenges in synchronizing model updates across heterogeneous edge devices.

 doi:10.20944/preprints202507.1162.v1

Elmahalwy and Mousa (2023) [14] propose a hybrid ensemble framework combining Isolation Forest and deep autoencoders for anomaly detection in decentralized systems. The model improves detection accuracy for complex multi-dimensional data, yet the computational cost of deep learning components limits its suitability for resource-constrained agent platforms.

Zhou et al. (2024) [15] introduce a lightweight decentralized identity framework for secure drone swarm communications. Their protocol leverages DIDs and verifiable claims for peer authentication, providing resilience against identity spoofing, though scalability for very large swarms remains an open research question.

Nguyen et al. (2023) [16] present a blockchain-integrated federated learning framework for multi-agent systems, aimed at enhancing trustworthiness in collaborative AI models. While the system ensures tamper-proof model aggregation, the added blockchain layer increases latency, raising concerns for time-sensitive applications.

Lin et al. (2024) [17] discuss Decentralized Physical Infrastructure Networks (DePIN) as a paradigm for agent-driven communication in physical systems. Their review outlines how secure protocols for AI agents can contribute to resilient and privacy-preserving physical infrastructures, but notes the lack of standardized security interfaces across platforms.

3. Methodology

3.1. Protocol Analysis

3.2. Framework Architecture

The proposed framework for secure communication among decentralized AI agents consists of five core components: **Decentralized Identity Module**, **Encrypted Communication Engine**, **AI-Inspired Threat Detection**, **Consensus Layer (Blockchain Integration)**, and **Resilience Module**.

Table 1. Comparison of Protocol Strengths and Limitations

ine **Protocol**	**Strengths**	**Limitations**
ine TLS/SSL	Provides robust encryption for client-server communications.	Relies on centralized certificate authorities, making it less suitable for decentralized systems.
ine Blockchain-based Protocols	Offers immutable logs and decentralized trust mechanisms.	Faces scalability issues, high latency, and potential privacy concerns due to transparent ledgers.
ine Decentralized Identifiers (DIDs)	Enables self-sovereign identity without centralized authorities.	Adoption is still emerging; lacks standardized implementation across platforms.
ine Agent-to-Agent (A2A) Protocols ine	Facilitates direct peer-to-peer communication with identity verification.	Early-stage development; lacks comprehensive threat detection mechanisms.

3.2.1. Decentralized Identity (DID) Module

Decentralized Identifiers (DIDs) and Verifiable Credentials (VCs) enable agents to establish and verify identities without relying on centralized authorities. This design prevents single points of failure common in traditional identity systems and mitigates risks such as impersonation and Sybil attacks. The DID module forms the foundation of trust, ensuring only authenticated agents can initiate or participate in communication.

3.2.2. Encrypted Communication Engine

The Noise Protocol Framework was selected for secure session key exchange, offering forward secrecy and post-compromise security. AES provides efficient and widely supported data encryption, while Post-Quantum Cryptography (PQC) options such as CRYSTALS-Kyber are incorporated for future resilience against quantum threats. The combination ensures confidentiality and integrity of messages while supporting adaptability to evolving cryptographic standards.

 doi:10.20944/preprints202507.1162.v1

3.2.3. AI-Inspired Threat Detection

While no model training or dataset evaluation was conducted in this study, the framework conceptually incorporates AI techniques—such as Isolation Forests for unsupervised anomaly detection and LSTM networks for analyzing sequential agent communication patterns. These models are cited in literature as promising for detecting unusual behaviors and emergent threats in decentralized environments. The inclusion reflects a design goal to enable dynamic, real-time detection without solely relying on static rules or signatures.

3.2.4. Consensus Layer (Blockchain Integration)

Blockchain is considered an optional component to provide immutable logging and auditability of critical security events (e.g., agent interactions, access attempts). It is suited for domains where transparency and regulatory compliance are priorities. However, due to its latency and scalability trade-offs, the blockchain layer is not mandatory and may not be applied in time-sensitive deployments like swarm robotics.

3.2.5. Resilience Module

The framework includes mechanisms for redundancy, backup, and rollback, ensuring graceful recovery from failures or attacks. This supports fault tolerance without introducing the complexity of fully decentralized state recovery systems, balancing robustness with operational simplicity.

3.2.6. Figure Interpretation

Figure 1 demonstrates how the five components of the framework interact sequentially and logically:

- Communication begins with **identity verification**.
- Data is **encrypted and securely transmitted**.
- The communication flow is **monitored for anomalies**.
- **Key events are optionally recorded** on the blockchain.
- The system's state is protected by the **resilience mechanisms**, ensuring reliability even under attack.

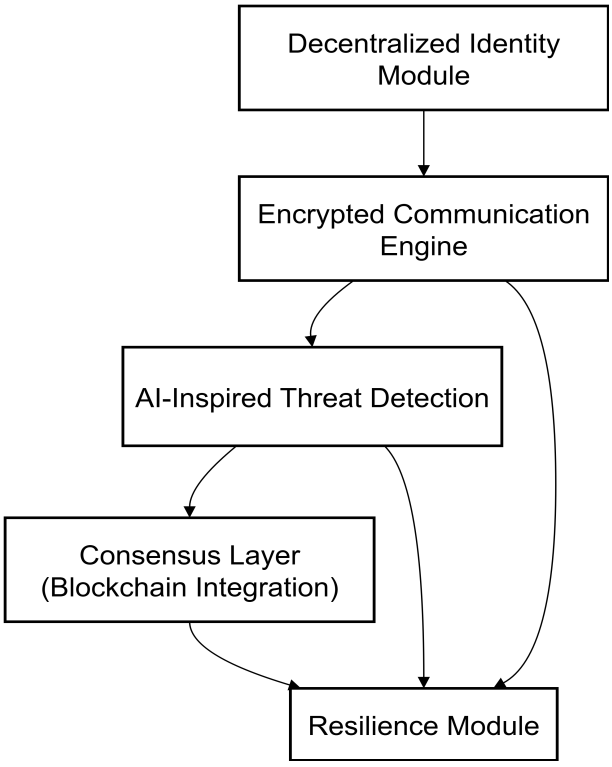

Figure 1. Framework architecture for secure communication among decentralized AI agents

 doi:10.20944/preprints202507.1162.v1

This layered design aligns with defense-in-depth principles, where multiple independent security controls work together to mitigate risk. The modular structure allows tailoring the architecture to specific domains—such as omitting blockchain for low-latency IoT applications or emphasizing AI detection in critical infrastructure deployments.

3.3. Operational Flow of the Proposed Framework

The diagram illustrates the dynamic interactions between decentralized AI agents and the key security components of the proposed framework. The process begins when Agent 1 initiates a secure session with Agent 2. This initial handshake involves encrypted data transmission and a challenge-response authentication mechanism to ensure mutual verification between the agents. Agent 2 then validates Agent 1's identity by querying a Decentralized Registry. The registry verifies the Decentralized Identifier (DID) and returns the corresponding verifiable credentials, allowing Agent 2 to confirm Agent 1's legitimacy without relying on a centralized certificate authority.

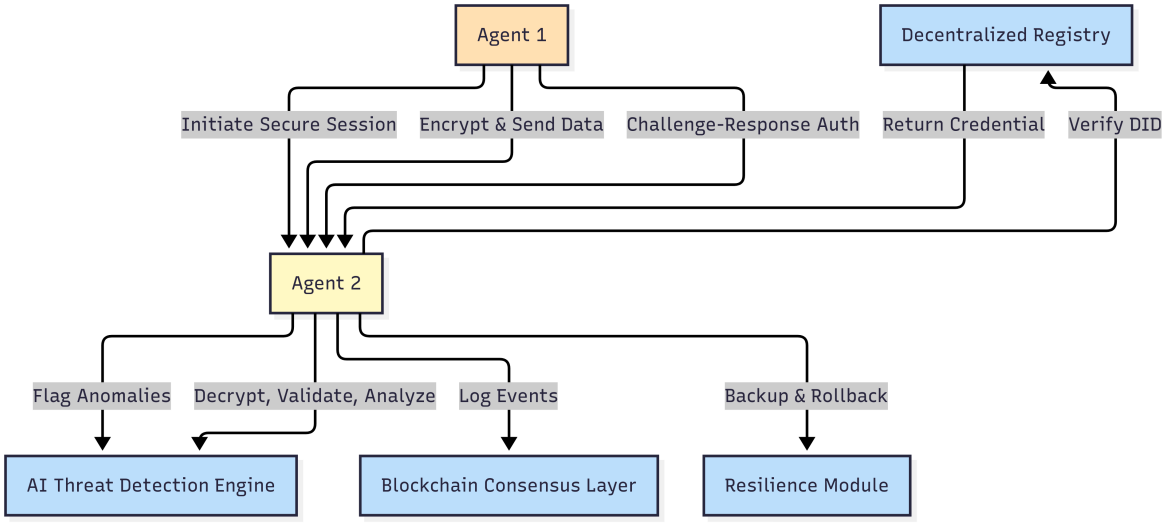

Figure 2. Working flow of secure communication between decentralized AI agents and key security components

Once the secure channel is established, Agent 2 decrypts, validates, and analyzes the incoming data through the Encrypted Communication Engine. This engine ensures that all data is protected from eavesdropping, tampering, or replay attacks. Concurrently, the AI Threat Detection Engine continuously monitors the communication flow, applying anomaly detection models to identify irregular patterns, potential intrusions, or compromised agent behavior. If anomalies are detected, relevant logs and events can optionally be recorded in the Blockchain Consensus Layer, creating an immutable audit trail that enhances transparency and accountability. This blockchain integration is particularly valuable in domains where regulatory compliance or forensic analysis is necessary.

The Resilience Module complements these security mechanisms by providing backup and rollback capabilities. It ensures that, in the event of a failure or attack, the system can restore a stable operational state without significant downtime or data loss. This module reinforces the system's fault tolerance and enables it to maintain service continuity even under adverse conditions.

Overall, the figure demonstrates how the framework combines multiple security layers—including decentralized identity verification, encryption, AI-powered monitoring, blockchain-based logging, and resilience mechanisms—to provide a holistic and adaptive security model for decentralized multi-agent systems. Each component works collaboratively to ensure secure, trustworthy, and robust communication. Importantly, this design eliminates reliance on centralized authorities, aligning with the principles of decentralization and self-sovereign agent governance. Moreover, the modular nature of the framework allows customization to different use cases; for example, blockchain logging can be enabled in regulatory-heavy environments, while lightweight deployments can skip it to reduce latency and overhead.

 doi:10.20944/preprints202507.1162.v1

3.3.1. Formal Protocol Specification and Sequence Flow

To provide a precise specification of the proposed secure communication process among decentralized AI agents, this section presents a formal protocol notation and sequence flow diagram. The model outlines key message exchanges, cryptographic operations, identity verification, anomaly detection, and optional audit logging.

Formal Protocol Flow

Let:

- A_1: Agent 1 (initiator)
- A_2: Agent 2 (responder)
- R: Decentralized identity registry
- SK: Session key
- $E_{AES}(M, SK)$: AES encryption of message M with key SK
- $\text{Sign}_{A_1}(N)$: Agent 1's signature of nonce N

The protocol proceeds as follows:

1. $A_1 \rightarrow A_2$: DID_{A_1}, session request
2. $A_2 \rightarrow R$: DID_{A_1}
3. $R \rightarrow A_2$: credentials_{A_1}
4. $A_2 \rightarrow A_1$: N (challenge nonce)
5. $A_1 \rightarrow A_2$: $\text{Sign}_{A_1}(N)$
6. $A_1 \leftrightarrow A_2$: Noise Protocol key exchange $\Rightarrow SK$
7. $A_1 \rightarrow A_2$: $E_{AES}(M, SK)$
8. $A_2 \rightarrow AIEngine$: $\text{meta}(M)$
9. $A_2 \rightarrow$ Blockchain (optional): log(event)

Sequence Diagram

The sequence diagram in Figure 3 illustrates these interactions in a visual format.

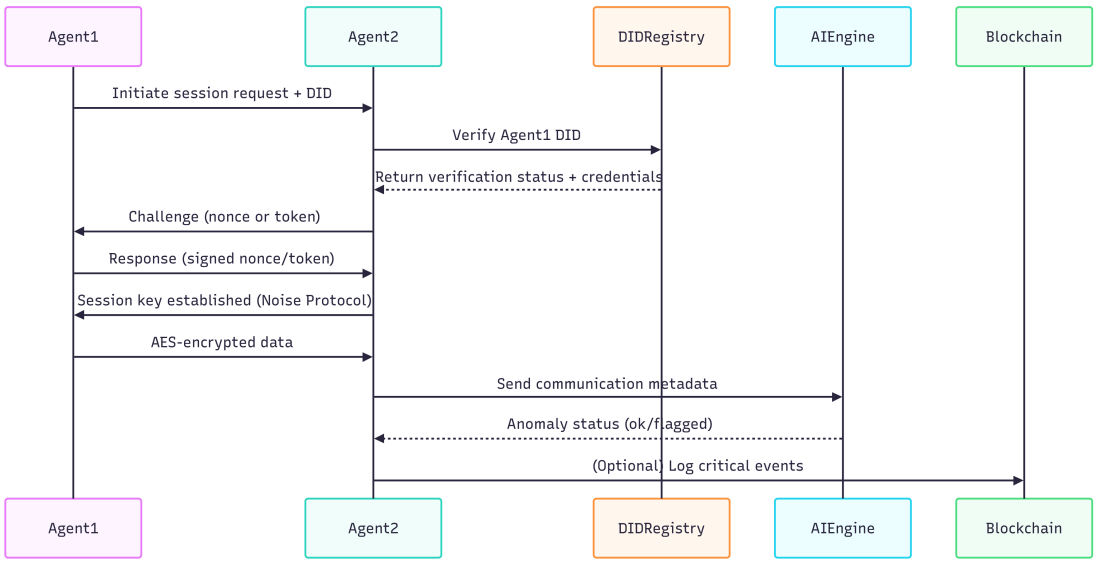

Figure 3. Sequence flow of secure communication between decentralized AI agents and framework components.

This formal model ensures clarity in the communication protocol, enabling future work on formal verification, simulation, or security proof construction.

 doi:10.20944/preprints202507.1162.v1

4. Results & Discussion

This section presents a conceptual analysis of the proposed framework for secure communication among decentralized AI agents. Since no dataset was used and no real-world implementation or simulation was conducted in this study, the discussion focuses on how the framework theoretically addresses known security challenges, synthesizing insights from related works and architectural considerations. The section highlights potential strengths, domain applicability, trade-offs, and areas for future validation.

4.1. Conceptual Strengths of the Framework

The framework integrates several key components that together provide a holistic security solution:

- **Decentralized Identity (DID) Module:** By leveraging Decentralized Identifiers (DIDs) and Verifiable Credentials (VCs), the framework enables autonomous, tamper-resistant identity verification. This eliminates dependence on centralized certificate authorities and mitigates impersonation, Sybil attacks, and unauthorized agent participation.
- **Encrypted Communication Engine:** The use of the Noise Protocol Framework for key exchange and AES combined with post-quantum options for data encryption ensures confidentiality, integrity, and forward secrecy. This design helps protect communication even in the presence of adversaries capable of future quantum decryption attacks.
- **AI-Inspired Threat Detection:** Although not implemented or trained in this study, the inclusion of AI techniques such as Isolation Forests and LSTM networks offers a path toward adaptive, real-time detection of anomalous communication patterns, enabling proactive threat mitigation.
- **Blockchain Integration (Optional):** The optional consensus layer allows for immutable logging and auditability of critical events. This is particularly beneficial in domains requiring forensic analysis, regulatory compliance, or high transparency.
- **Resilience Module:** The system includes design provisions for backup, rollback, and failover, supporting continuous operation even in the face of security incidents or system faults.

These components collectively provide defense-in-depth, aligning with best practices in decentralized system security.

4.2. Use Cases and Practical Implications

This framework is conceptually designed to address real-world challenges in securing decentralized AI agent communication across diverse domains. Its modular architecture allows adaptation to specific industries, aiming to provide both security and scalability. Below, we present a technical overview of potential applications and a conceptual scenario illustrating its use in a drone delivery network.

4.3. Conceptual Overview by Domain

4.4. Conceptual Scenario: Drone Delivery Network

In a decentralized drone delivery system operating across a city, the framework could conceptually function as follows:

1. **Drone Launch:** Each drone initializes its identity using a DID issued by a decentralized registry.
2. **Pre-Flight Communication:** Drones exchange flight plans and cargo details; messages are encrypted with AES, with session keys negotiated via the Noise Protocol.
3. **Identity Verification:** Upon receiving a message, a drone validates the sender's DID and verifiable credential; invalid messages are discarded.
4. **Ongoing Monitoring:** Drone behaviors such as speed, altitude, and route adherence are monitored. The AI Threat Detection Engine conceptually analyzes these logs to flag anomalies.
5. **Anomaly Detected:** Suspicious behavior (e.g., sudden altitude drop) triggers an automated alert and avoidance instructions to nearby drones.

 doi:10.20944/preprints202507.1162.v1

6. **Blockchain Logging (Optional):** Critical events (e.g., identity verifications, alerts) can be logged to a blockchain for auditability.
7. **Recovery:** If an attack is detected, the Resilience Module enables rollback of affected states and adjusts delivery schedules.
8. **Post-Operation Analysis:** Collected data could inform future model refinements, although no actual retraining or testing was performed in this study.

Table 2. Conceptual application of the framework across domains

ine **Domain**	**Implementation Highlights**	**Framework Application**
ine Smart Grids	Conceptually, agents would manage distributed energy resources, verify peers via DIDs, encrypt load data using AES, and apply AI-based anomaly detection to identify events like sudden demand spikes.	Prevents data injection, supports load balancing, and helps maintain grid stability.
ine Autonomous Logistics (Drones)	Drones would exchange location and status data securely. DIDs authenticate each drone; AES encrypts communications; anomaly detection identifies potential route deviations.	Protects route sharing, prevents hijacking, ensures mission reliability.
ine Healthcare IoT	Devices such as wearables and monitors would verify identities via DIDs, encrypt patient data in transit, and monitor for unusual patterns in vital readings.	Enhances patient confidentiality, prevents data leaks, improves medical data accuracy.
ine Financial Systems	AI agents would validate transactions using DIDs, encrypt data transmissions, and detect fraud attempts (e.g., anomalous fund transfers) using AI anomaly detection.	Secures transactions, prevents fraud, ensures regulatory compliance.
ine		

4.5. Trade-offs, Design Considerations, and Scalability

The framework's layered approach brings significant conceptual security advantages, but it also introduces potential challenges:

- **Performance Overhead:** Integrating advanced cryptography (e.g., post-quantum algorithms) and optional blockchain logging may introduce latency and resource usage concerns, particularly in environments like swarm robotics or low-power IoT where minimal delay is critical.
- **Complexity of Deployment:** The combination of decentralized identity systems, cryptography, AI models, and optional blockchain adds deployment complexity. Real-world implementations would require careful orchestration, configuration, and maintenance.

 doi:10.20944/preprints202507.1162.v1

- **Scalability Considerations:** As the number of agents grows, ensuring efficient key management, anomaly detection responsiveness, and blockchain transaction handling will be essential. Without empirical validation, the scalability of this architecture remains theoretical.

4.6. Future Work and Limitations

This study presents a conceptual framework for securing communication among decentralized AI agents, but several limitations must be acknowledged. No dataset or prototype implementation was applied, so the framework's performance, latency characteristics, scalability, and threat detection accuracy remain untested. Quantitative benchmarks cannot yet be reported, and challenges such as packet loss, unpredictable network latency, or sophisticated adversary behavior have not been addressed in practice.

Future work will focus on developing and empirically validating the framework through:

- **Prototype development:** Building proof-of-concept implementations using Python (e.g., FastAPI for agent simulation), DIDKit for identity, PyCryptodome for AES encryption, and Noise Protocol bindings for key exchange.
- **Testbed deployment:** Evaluating in specific environments, including:
  - *Smart grids:* Using GridLAB-D or IEEE PES test systems to assess agent coordination, latency, and detection precision during load balancing tasks.
  - *Drone logistics:* Simulating urban drone networks in environments like NASA's AAM ecosystem or FAA corridors to measure communication delays, anomaly detection accuracy, and resilience.
  - *Healthcare IoT:* Testing secure device communication in lab environments, with attention to privacy preservation and data integrity.
  - *Financial systems:* Using cyber-range platforms to validate secure transaction flows and fraud detection performance.
- **Formal validation:** Applying model checking (e.g., TLA+, ProVerif) to verify correctness, safety, and security properties of the protocol.
- **Lightweight adaptation:** Exploring reduced-overhead configurations for resource-constrained environments where components like blockchain logging may not be feasible.
- **Interoperability analysis:** Assessing integration with existing agent communication standards and platforms.

Finally, future work must consider ethical, privacy, and legal dimensions. This includes evaluating potential risks from false positives in anomaly detection, managing privacy concerns with immutable blockchain logs under data protection regulations (e.g., GDPR, CCPA), and clarifying the legal accountability of autonomous agents operating across jurisdictions. Addressing these dimensions will be essential for responsible and compliant deployment of the framework.

4.7. Discussion

The proposed framework offers a layered security architecture for decentralized AI agent communication, conceptually addressing key challenges in identity verification, data confidentiality, integrity, threat detection, and system resilience. By integrating decentralized identifiers, hybrid cryptographic techniques, AI-inspired anomaly detection, and optional blockchain-based logging, it aligns with best practices for defense-in-depth in decentralized systems.

A primary advantage is its modularity, which allows adaptation to domain-specific requirements without reliance on centralized authorities. This design supports self-sovereign authentication, proactive threat monitoring, and transparent auditability. However, the use of advanced cryptographic methods, AI components, and optional blockchain integration may increase latency, resource consumption, and implementation complexity, particularly in resource-constrained or ultra-low-latency environments.

 doi:10.20944/preprints202507.1162.v1

Future work should focus on empirical validation through prototype development and domain-specific testbed deployments. This would enable quantitative assessment of trade-offs between security, performance, and scalability, and guide refinements for practical adoption.

5. Conclusions

This paper presents a layered security framework for decentralized AI agent communication, integrating decentralized identity verification, hybrid encryption, AI-inspired anomaly detection, and optional blockchain audit logging into a cohesive architecture. The framework's modular design enables domain-specific customization, supporting secure communication across diverse applications such as smart grids, autonomous logistics, healthcare IoT, and financial systems. A formal protocol specification provides a foundation for future verification and implementation. Future work will focus on prototype development, empirical validation in domain-relevant testbeds, and addressing ethical, privacy, and legal considerations in deployment.

This work highlights the importance of holistic, modular security architectures in enabling trustworthy decentralized AI systems.

References

1. Huang K., Narajala V. S., Yeoh J., Raskar R., Harkati Y., Huang J., Habler I. & Hughes C. A Novel Zero-Trust Identity Framework for Agentic AI: Decentralized Authentication and Fine-Grained Access Control. arXiv preprint arXiv:2505.19301 (2025).
2. Ehtesham A., Singh A., Gupta G. K. & Kumar S. A Survey of Agent Interoperability Protocols: MCP, ACP, A2A, and ANP. arXiv preprint arXiv:2505.02279 (2025).
3. Ranjan R., Gupta S. & Singh S. N. LOKA Protocol: A Decentralized Framework for Trustworthy and Ethical AI Agent Ecosystems. arXiv preprint arXiv:2504.10915 (2025).
4. Wang Q., Zhang Y. & Li H. AI Agents Meet Blockchain: A Survey on Secure and Scalable Collaboration for Multi-Agents. Future Internet 17, 57 (2025).
5. Nagothu D., Pentapati A., Sabharinadh R., Gumpula R., Uppu S. K. & Sharma V. Lightweight Consensus Algorithms for IoT Agent Networks. Int. J. Comput. Sci. 12, 112–118 (2022).
6. Liu Y., Park J. H. & Shin K. LSTM-based Anomaly Detector for Swarm Systems. IEEE Internet Things J. 10, 5124–5132 (2023).
7. Chalvatzakis T. et al. Unsupervised Anomaly Detection with Isolation Forests in Industrial Agent Networks. IEEE Access 12, 9981–9995 (2024).
8. Sharma S. et al. Blockchain-Logged DID Protocol for IIoT Security. Int. Conf. IoT Security 2024, 144–150 (2024).
9. IJCAI 2023 Participants. Cooperative Anomaly Detection Scheme for Decentralized MAS. IJCAI 2023 Proceedings, 3421–3427 (2023).
10. Andreu A. A Survey on Reputation Systems and Consensus for Decentralized Agentic AI. arXiv preprint arXiv:2503.14100 (2025).
11. Li X., Zhao H., Chen Y. & Wang S. SAFEFLOW: Protocol-Level Info-Flow Control for LLM-Empowered Agents. arXiv preprint arXiv:2502.09456 (2025).
12. Hernandez-Ramos J. L. et al. Agent-Centric Security Architecture with DIDs. IEEE Access 11, 21345–21359 (2023).
13. Pavle M., Matic M. & Popovic M. Federated Isolation Forest for Edge-Based IoT Systems. arXiv preprint arXiv:2506.05138 (2025).
14. Elmahalwy A. & Mousa H. Hybrid Ensemble Framework for Decentralized Anomaly Detection. Int. J. Electr. Comput. Eng. 13, 82–91 (2023).
15. Zhou L., Li Z., Chen R. & Wang X. Decentralized Identity Framework for Drone Swarm Communications. IEEE Commun. Mag. 59, 86–92 (2024).
16. Nguyen N. T., Albalushi A. S. & McDonald A. Blockchain-Integrated Federated Learning for Multi-Agent Systems. J. Netw. Comput. Appl. 2023, 1–15 (2023).
17. Lin Y., Wang T., Shi L. & Zhang S. Decentralized Physical Infrastructure Networks (DePIN): Challenges and Opportunities. IEEE Netw. 38, 58–66 (2024).

 doi:10.20944/preprints202507.1162.v1