Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ **Research Article**

Autonomous Cybersecurity and Data Governance in Distributed Financial Systems: A Multi-Agent AI Framework for Threat Detection, Policy Enforcement, and Incident Response

David Excel Ozowara (0009-0003-2386-2949)[1*], Chukwudera Obumneke Anunagba (0009-0003-3949-5290) [2]

[1] *Goldman Sachs, Dallas, Texas, USA.*

[2] *Barclays, Whippany, New Jersey, USA.*

***Corresponding author:** David Excel Ozowara, daveozzy6@gmail.com*

***Author emails:** David Excel Ozowara: daveozzy6@gmail.com; Chukwudera Obumneke Anunagba: chiokemuo@gmail.com*

ARTICLE INFO

Received: 03 Sept 2025

Revised: 17 Oct 2025

Accepted: 28 Oct 2025

ABSTRACT

Distributed financial systems integrate banking platforms, cloud services, application programming interfaces, transaction processors, and data repositories, creating interconnected attack surfaces that challenge conventional centralized cybersecurity mechanisms. This review synthesizes research on multi-agent artificial intelligence (AI), collaborative intrusion detection, federated threat intelligence, reinforcement learning, data governance, and autonomous incident response. It identifies a persistent architectural gap between detection models, enforceable institutional policies, and validated recovery actions. To address this gap, the paper proposes a governance-aware multi-agent reference architecture that combines distributed sensing, provenance-preserving evidence fusion, data-quality-aware risk scoring, run-time policy authorization, hierarchical response coordination, and human oversight. A formal decision model represents collaborative threat assessment and constrained action selection while accounting for operational disruption, uncertainty, and regulatory obligations. The paper further specifies a comparative evaluation protocol spanning detection effectiveness, response latency, unauthorized action rates, service continuity, privacy exposure, and audit completeness. The synthesis indicates that greater autonomy is defensible only when its authority is bounded, its decisions are traceable, and its actions can be reviewed or reversed. The proposed framework provides a research blueprint rather than a claim of demonstrated production deployment.

Keywords: Multi-agent artificial intelligence; financial cybersecurity; data governance; distributed intrusion detection; autonomous incident response; policy enforcement; federated learning; trustworthy AI.

1 INTRODUCTION: THE CASE FOR GOVERNED AUTONOMY IN DISTRIBUTED FINANCIAL SECURITY

1.1 Distributed financial infrastructures as adaptive cyber-physical environments

Financial infrastructures now combine legacy banking systems, cloud services, APIs, microservices, algorithmic trading platforms, data lakes, and connected operational technologies. This integration increases processing capacity while coupling cyber events to transaction integrity, service availability, and, in some settings, physical processes. The resulting environment changes continuously as users, applications, devices, and adversaries interact. Research on financial cyber-physical systems therefore frames security as a joint problem of anomaly detection, threat modeling, data governance, and response coordination rather than as a perimeter-control problem. Cyber-physical studies similarly show that physical faults and false-data-injection attacks can produce

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

related observable anomalies, requiring detectors that distinguish malicious behavior from operational disturbance (Li et al., 2020).

1.2 From isolated security tools to coordinated multi-agent defense

Conventional signatures, static rules, and isolated monitoring tools struggle when attacks change across hosts, services, and organizational boundaries (Finistrella et al., 2025). **Multi-agent systems** distribute observation and decision-making across specialized components that can exchange partial evidence, maintain local autonomy, and coordinate actions. This architecture supports adaptive detection, policy enforcement, investigation, and mitigation while reducing dependence on a single control point. The shift from task-specific agents to agentic systems adds dynamic decomposition, persistent memory, tool use, and coordinated autonomy, but also introduces coordination failure and emergent behavior (Li et al., 2024).

1.3 Scope, analytical focus, and organizing framework of the review

This review examines governed autonomy in distributed financial security through six linked functions: architecture, detection, intelligence sharing, policy execution, incident response, and trust controls. It treats reinforcement learning, federated learning, neuro-symbolic reasoning, and large language models as complementary mechanisms rather than interchangeable solutions. The analysis distinguishes benchmark performance from operational validity, since most reported evaluations use simulations or public datasets rather than live regulated environments. The central claim is conditional: autonomy can reduce detection and response delays only when agents operate within explicit policy boundaries, secure communication channels, auditable workflows, and human escalation paths.

1.4 Review Approach and Analytical Classification

This article adopts a structured narrative synthesis rather than a systematic review or meta-analysis. The cited literature is organized around six interconnected analytical dimensions: distributed agent architecture, threat detection, intelligence sharing, governance and policy execution, incident response, and trustworthy autonomy. Studies are differentiated by contribution type, including surveys and taxonomies, conceptual architectures, algorithmic proposals, benchmark evaluations, simulation-based demonstrations, and operational case studies. This classification prevents design proposals from being interpreted as evidence of deployed effectiveness.

Cross-study interpretation prioritizes the task definition, data source, evaluation environment, comparison baseline, security and operational metrics, and limitations reported by each publication. Reported accuracy and latency values are not pooled because datasets, threat models, sample sizes, and measurement procedures are heterogeneous. The synthesis therefore identifies recurring architectural mechanisms and evidence gaps rather than estimating a single numerical effectiveness measure.

2 FROM CENTRALIZED CONTROLS TO DISTRIBUTED AGENTIC CYBERSECURITY ARCHITECTURES

2.1 Architectural principles of multi-agent coordination across heterogeneous infrastructure

Distributed agentic security begins by assigning bounded responsibilities to agents positioned near relevant data and control surfaces. Host, network, identity, application, transaction, and infrastructure agents can produce local assessments while a coordination layer reconciles evidence and allocates tasks. This arrangement addresses the scale and heterogeneity that make centralized analysis slow or vulnerable to a single point of failure. A multi-agent IDS using distributed deep reinforcement learning describes this pattern as a fault-tolerant, multi-view detection platform in which local agents coordinate classification across network observations (Sethi et al., 2021). Big-IDS similarly combines decentralized cooperative learning, protected communication, concurrent preprocessing, cloud resources, and streaming analysis for large data flows (Louati et al., 2024).

Coordination requires more than message exchange. Agents need role definitions, state representations, communication protocols, action constraints, and mechanisms for resolving disagreement. The Major-Minor-RL model assigns a principal classifier to predict traffic status while auxiliary agents correct that decision when collective evidence diverges (Shi & He, 2021). Hierarchical organization offers another solution: higher-level agents

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

set objectives or allocate resources, while lower-level agents execute host, service, or network actions. Tang and colleagues formulate this structure through a Stackelberg hypergame and hierarchical reinforcement learning, using role separation to reduce action-space complexity under incomplete information (Tang et al., 2024). Organizational constraints can also be represented symbolically and imposed during policy learning (Li et al., 2024), linking learned behavior to designer-defined roles and objectives.

Heterogeneous financial infrastructure makes this separation especially relevant. A cloud workload agent may isolate a container, while a transaction agent assesses fraud exposure and a compliance agent evaluates whether a proposed action is permissible. These agents should share normalized evidence without collapsing their distinct decision criteria. Secure, ontology-aligned messages and mappings to common threat descriptions can support this exchange; MALCDF uses encrypted communication and MITRE ATT&CK-aligned outputs to make multi-agent results more consistent and reviewable. The architecture must also preserve local failure containment, because an erroneous agent should not automatically obtain unrestricted authority over payment, identity, or production systems.

2.2 Proposed Governance-Aware Multi-Agent Reference Architecture

Building on the architectural patterns identified in the literature, this paper proposes a governance-aware reference architecture organized into six interacting layers. The observation layer consists of host, network, identity, application, transaction, and infrastructure agents that collect timestamped and source-identified evidence. The analytical layer combines local anomaly estimates with fraud indicators, contextual intelligence, and uncertainty measurements. A coordination layer reconciles cross-domain observations, identifies potentially related events, and maintains a shared incident state without requiring wholesale transfer of sensitive raw data. The governance layer validates agent identity, delegated authority, data classification, policy version, and the conditions under which an action is permitted. The response layer selects among monitoring, investigation, reversible containment, escalation, and recovery actions. Finally, an assurance layer records decisions, evaluates execution outcomes, detects model drift, and supports post-incident review.

Unlike an unconstrained autonomous workflow, this architecture places an authorization gate between action recommendation and execution. Each proposed intervention must be attributable to an agent, linked to supporting evidence, checked against current policy, and evaluated for its potential effect on dependent financial services. The coordination layer may recommend a transaction hold or workload isolation; it cannot infer authority merely from a high anomaly score. High-impact or irreversible actions require designated approval, whereas predefined low-impact and reversible interventions may execute automatically within explicit limits. This separation makes local responsiveness compatible with institution-wide accountability.

Let the distributed environment at time t be modeled as Equation (1), where the environment comprises protected assets, dependencies, permissible actions, and versioned institutional policies.

Each observation agent generates an evidence item according to Equation (2). This representation incorporates local threat scores, timestamps, data-quality information, confidence estimates, and provenance identifiers. These fields provide a common representation for subsequent evidence fusion and audit reconstruction.

Equation (1): Distributed multi-agent environment

$$\boxed{\mathcal{G}_t = (\mathcal{V}_t, \mathcal{E}_t, \mathcal{A}, \mathcal{P}_t)}$$

Where:

- $\mathcal{G}_t$ = distributed financial cybersecurity environment at time \(t\).
- $\mathcal{V}_t$ = protected assets and agents.
- $\mathcal{E}_t$ = dependencies and evidence-sharing relationships.
- $\mathcal{A}$ = set of permissible action types.
- $\mathcal{P}_t$ = versioned institutional policy state.

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ **Research Article**

Equation (2): Agent evidence representation

$$e_{i,t} = (s_{i,t}, \tau_{i,t}, q_{i,t}, c_{i,t}, \pi_{i,t})$$

Where:

- $e_{i,t}$ = evidence generated by agent \(i\) at time \(t\).
- $s_{i,t}$ = local threat score.
- $\tau_{i,t}$ = observation timestamp.
- $q_{i,t}$ = data-quality indicator.
- $c_{i,t}$ = agent confidence.
- $\pi_{i,t}$ = evidence provenance identifier.

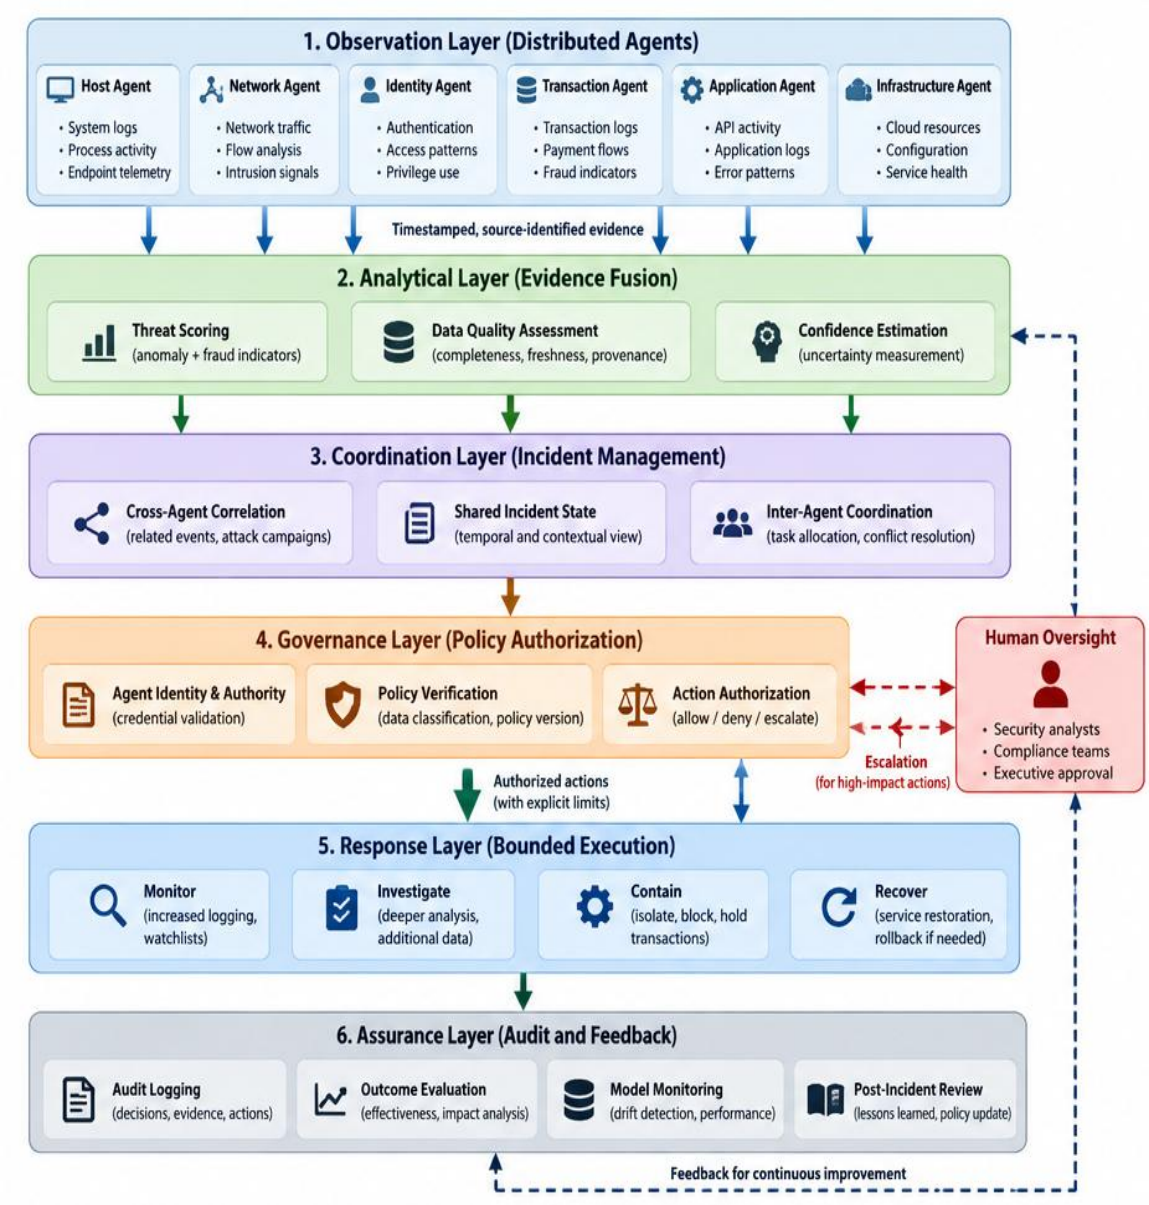


Figure 1. Governance-aware multi-agent cybersecurity reference architecture for distributed financial systems.

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

The proposed architecture integrates distributed observation agents, collaborative threat analysis, incident coordination, policy authorization, bounded response execution, and continuous assurance. The framework incorporates human oversight and feedback mechanisms to support accountable autonomous cybersecurity operations.

Figure 1 illustrates the proposed six-layer governance-aware multi-agent cybersecurity architecture designed to support autonomous threat detection, policy enforcement, and incident response across distributed financial infrastructures. The observation layer comprises specialized host, network, identity, transaction, application, and infrastructure agents responsible for collecting security telemetry and identifying potentially anomalous behavior. The analytical layer combines threat scoring, data-quality assessment, and confidence estimation to produce contextual security assessments. The coordination layer correlates observations from multiple agents, maintains shared incident states, and facilitates collaborative decision-making.

The governance layer serves as a mandatory authorization boundary, validating agent identities, institutional policies, delegated permissions, and action restrictions before any intervention is executed. Authorized responses are subsequently implemented through the bounded response layer, which supports monitoring, investigation, containment, and recovery. Finally, the assurance layer records operational activities, evaluates intervention outcomes, monitors model behavior, and supports post-incident analysis. Human oversight provides additional authorization and escalation pathways for high-impact security decisions. The architecture emphasizes the separation between intelligent threat assessment and permissible operational execution, ensuring that autonomous security interventions remain traceable, proportionate, and subject to institutional governance.

2.3 Reinforcement learning, hierarchical decision-making, and autonomous defense

Reinforcement learning frames defense as sequential decision-making under changing observations, actions, and rewards. This is useful where an agent must choose among monitoring, investigation, access restriction, deception, containment, or recovery actions rather than merely assigning an attack label. Surveys classify RL approaches by both agent organization and deployment setting, including host-based, network-based, and software-defined networking environments (Finistrella et al., 2025). MARL extends the formulation by allowing agents to learn cooperative or adversarial strategies, but it also introduces non-stationarity, communication costs, computational complexity, and dependence on reward design (Landolt et al., 2025).

Hierarchical methods reduce exploration difficulty by separating strategic decisions from local execution. In CybORG simulation, hierarchical MARL generated coordinated defense strategies and was evaluated against alternative approaches for strategy generation and control (Tang et al., 2024). Grouped MARL applies a related principle to cloud defense by allowing defenders within groups to share information and experience, reducing state-space explosion while balancing defense rewards against resource and cost constraints (Chen et al., 2025). These results support decomposition as an engineering response to scale, but simulation success does not establish safe behavior under production constraints, distribution shift, or adversarial manipulation.

Autonomous defense also depends on the quality and stability of the environment model. CAGE Challenge 4 was designed to test MARL in enterprise-network defense because single-agent approaches become difficult to scale across large networks (Kiely et al., 2025). Transferable anomaly detection research addresses a different limitation by learning from interactive processes and prioritizing outliers without requiring complete relabeling of each new dataset (He et al., 2024). The common implication is that agents should adapt to changing traffic and attack behavior, while retaining constraints that prevent exploratory actions from damaging financial services.

2.4 Financial-sector applications across cloud, microservices, trade finance, and cyber-physical systems

Applications span infrastructure operations and business workflows. In hybrid-cloud financial systems, a performance agent can optimize database activity while a governance agent applies encryption and personally identifiable information obfuscation, linking resource orchestration to compliance controls (Sanepalli, 2025). Agentic frameworks for financial infrastructure (Li et al., 2024) similarly assign specialist capabilities across security response, database optimization, capacity management, and network performance, with self-restoration

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

intended to address failures before service loss. These designs treat security as part of service operation rather than as a separate SOC function.

Trade finance illustrates the need to connect agent autonomy with domain rules. TradeQu (Li et al., 2024) places a versioned, tenant-scoped Policy Store, a graph of entities and documents, and an orchestrator with policy gates at the center of document verification and compliance execution. The architecture is materially different from a generic alert pipeline because the agent may recommend or execute actions affecting letters of credit, sanctions screening, customer records, and transaction release. Financial-services reviews identify related applications in risk management, trading, compliance, fraud detection, and smart-contract orchestration, but also report limited long-term empirical evidence and unresolved tensions among scalability, regulation, and real-time performance (Adeliyi, 2025).

Cloud and cyber-physical applications add operational consequences. Grouped MARL addresses competing attackers and defenders allocating limited cloud resources (Chen et al., 2025), while an autonomous response controller for cyber-physical systems combines a hierarchical risk tree with a competitive Markov decision process to select mitigation actions. In a generator testbed, that controller responded to an Aurora attack and recovered the system faster than the compared static and Suricata-based responses (Kholidy, 2020). The architectural lesson is that distributed agents require explicit authority boundaries, because an action that is reversible in a cloud workload may be unsafe when applied to a payment ledger, trading engine, or physical control system.

3 ADAPTIVE THREAT DETECTION, ANOMALY IDENTIFICATION, AND DISTRIBUTED INTELLIGENCE SHARING

3.1 Collaborative intrusion detection for known, unknown, and coordinated attacks

Detection systems differ in whether they recognize known signatures, infer deviations from behavior, or combine both forms of evidence. Hybrid NIDS research combines a signature-based component with anomaly detection and machine-learning classifiers to address known and zero-day attacks (Ouiazzane et al., 2021). The design is attractive for financial environments because signatures can support precise recognition of established malware or exploit patterns, while anomaly models can identify unusual account, service, or network behavior. Yet benchmark accuracy alone does not establish robustness against concept drift, class imbalance, or deliberate manipulation of normality baselines.

Distributed agents address these limitations by observing different traffic segments and exchanging learned representations. Soltani and colleagues propose continual anomaly detectors at local sensors, federated sharing of knowledge, and sequential packet labeling that updates an attack probability as additional packets arrive (Soltani et al., 2024). Their evaluation on CIC-IDS2017 and CSE-CIC-IDS2018 reports that CNN-based models adapted to traffic drift with an average detection rate above 95 percent after an updating phase using 128 new flows, while LSTM models detected intrusions after observing the first 15 packets (Soltani et al., 2024). These endpoints concern the reported datasets and model settings, not general production performance. They nevertheless show why sequential and distributed inference can be preferable to waiting for complete flows.

Collaborative RL adds decision adaptation to distributed observation. Attention-based multi-agent IDS research uses multiple DQN agents to construct a coordinated, multi-view detector and evaluates robustness against a black-box adversarial attack, with a denoising autoencoder used as an additional defense (Sethi et al., 2021). Big-IDS reports 97.44 percent accuracy on NSL-KDD while using decentralized MARL, cloud processing, and streaming analysis (Louati et al., 2024). The two studies differ in dataset, architecture, and evaluation setting, so their percentages should not be compared as a common league table. Their shared contribution is architectural: local sensing becomes more useful when agents combine partial signals without requiring all raw traffic to move to a central analyzer.

Financial fraud detection presents a parallel problem with different data semantics. A review of 47 studies reports detection rates of 87 to 94 percent and false-positive reductions of 40 to 60 percent for contemporary AI systems compared with traditional rule-based methods, while integrated supervised and unsupervised approaches were more effective for novel fraud patterns (Olowu et al., 2024). These findings support layered detection, but they do

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

not show that a multi-agent architecture caused the reported improvements. In financial systems, the detector must therefore preserve the distinction between a cyber intrusion, a fraudulent transaction, an unusual but legitimate customer action, and a data-quality anomaly.

Coordinated attacks require agents to reason across time and domains. A local sensor may observe credential misuse, while an identity agent sees unusual privilege escalation and a transaction agent detects abnormal payment routing. Cooperative detectors for multi-agent systems have used independent and graph-informed detectors to distinguish physical faults from false-data injection (Li et al., 2020). The same principle can be adapted to financial infrastructure, provided the shared representation records provenance, confidence, temporal context, and the limits of each observation.

Equation (3): Collaborative threat-risk assessment

$$R_{j,t} = \frac{\sum_{i=1}^{N} w_i \, q_{i,t} c_{i,t} s_{i,j,t}}{\sum_{i=1}^{N} w_i \, q_{i,t} c_{i,t}}$$

Where:

- $R_{j,t}$ = aggregated threat score for incident \(j\) at time \(t\).
- N = number of participating agents.
- w_i = governance-approved relative weight of agent \(i\).
- $q_{i,t} \in [0,1]$= evidence quality.
- $c_{i,t} \in [0,1]$= calibrated source confidence.
- $s_{i,j,t} \in [0,1]$= agent threat score.
- $w_i \geq 0$= nonnegative weighting constraint.

Equation (4): Evidence freshness decay

$$f_{i,t} = \exp\left[-\lambda_i \left(t - \tau_{i,t}\right)\right]$$

Subject to:

$$t \geq \tau_{i,t}, \qquad \lambda_i \geq 0$$

Where:

- $f_{i,t}$= freshness of the evidence from agent \(i\).
- λ_i= source-specific decay parameter.
- $t - \tau_{i,t}$= elapsed time since the evidence was observed.

A higher decay parameter causes the influence of older observations to decrease more rapidly. The purpose is to prevent stale evidence from disproportionately influencing real-time response decisions.

Technical limitation: The weighted score assumes a useful degree of reliability in the individual source assessments. It does not independently solve correlated-evidence bias or malicious-agent manipulation.

3.2 Data-Quality-Aware Collaborative Threat Assessment

Collaborative detection depends not only on the number of agents reporting suspicious activity but also on the reliability, timeliness, and independence of their observations. Financial telemetry may be incomplete because of network delays, intermittent services, failed logging, inconsistent transaction identifiers, and differences in institutional schemas. If these limitations are ignored, an agent ensemble may convert multiple low-quality observations into unjustifiably high confidence. The proposed framework therefore separates a source's estimated threat probability from the quality and credibility of its supporting evidence.

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

For an event or incident evaluated at a given time, a normalized collaborative risk estimate is defined in Equation (3). This formulation combines individual threat assessments using governance-approved source weights, data-quality indicators, and confidence estimates.

Equation (3) is evaluated only when the denominator is positive; otherwise, the system records an insufficient-evidence state and escalates rather than treating the event as benign. Evidence-quality scores may incorporate completeness, freshness, schema validity, and provenance integrity. This weighted aggregation is a transparent baseline rather than a complete probabilistic fusion rule: correlated signals, compromised participants, and distribution shifts require dedicated detection and robustness testing.

To account for observation age, an optional freshness component is introduced in Equation (4). This component can contribute to the evidence-quality score after timestamps are normalized to a common time unit. A delayed observation consequently loses influence at a rate reflecting the source and use case, while remaining available for forensic reconstruction. The weights and decay constants must be tuned and audited on institution-appropriate validation data rather than assumed to be universally applicable.

3.3 LLM-based anomaly interpretation, threat hunting, and alert verification

Large language models add interpretive capacity to detection pipelines by relating structured alerts to incident reports, threat knowledge, procedures, and analyst questions. Their strongest use in the reviewed literature is not autonomous classification of raw telemetry, but coordination of specialized functions. An LLM-based financial-market framework assigns agents to data conversion, expert analysis, institutional knowledge retrieval, cross-checking, and report consolidation to verify anomalies in S&P 500 data (Park, 2024). This structure addresses alert interpretation and human verification rather than replacing statistical detection.

Cybersecurity implementations use a similar division of labor. MALCDF separates Detection, Intelligence, Response, and Analysis agents, with secure communication and ATT&CK mappings; on a 50-record stream derived from the CICIDS2017 schema, it reports 90.0 percent detection accuracy, an 85.7 percent F1-score, a 9.1 percent false-positive rate, and 6.8 seconds of average per-event latency. The small stream and narrow evaluation constrain generalization. An enterprise threat-hunting framework reports processing more than 8,000 security events in 2.2 seconds, a 28 percent anomaly detection rate, complete pipeline success, and more than 205 SOAR workflows per mission (Swarna, 2025). Because the anomaly rate is not equivalent to precision or recall, the result should be interpreted as a pipeline-performance report rather than as proof of superior detection quality.

LLMs can also support investigative decomposition. Agents may retrieve relevant telemetry, query external or internal knowledge, compare competing explanations, and draft an evidence-linked report. AutoBnB evaluates centralized, decentralized, and hybrid LLM teams in the Backdoors & Breaches tabletop framework, finding that multi-agent collaboration can improve decision-making and adaptability in simulated incident response (Liu, 2025). Such evaluations test workflow coordination under controlled scenarios. They do not resolve whether an LLM agent can safely execute a high-impact action when telemetry is incomplete, conflicting, or adversarial.

The principal limitations concern grounding and tool control. Reviews identify prompt injection, insecure output handling, data poisoning, adversarial instructions, and denial-of-service risks in generative cybersecurity systems (Ferrag et al., 2025). A lifecycle governance model maps separate controls to perception, decision, memory, and execution, reflecting the possibility that an agent can ingest false evidence, hallucinate a rationale, retain poisoned information, or issue a harmful command. LLMs should therefore interpret and prioritize evidence within schema-bound tools, retrieval controls, output validation, and human escalation, rather than receive unrestricted access to financial systems.

3.4 Decentralized threat intelligence and privacy-preserving learning across institutions

Threat intelligence sharing creates a tension between collective visibility and institutional confidentiality. Centralized repositories can impose latency, concentration risk, and dependence on a single intermediary. A distributed cyber-intelligence architecture (Alaeifar et al., 2024). assigns local agents to anomaly detection, adversarial-behavior prediction, and threat scoring, then combines knowledge through federated learning and blockchain-supported sharing. The reported evaluation describes higher detection accuracy, lower response

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

latency, and fewer false positives than a centralized comparison, but the reported evaluation does not establish how these metrics were measured across institutions or how governance obligations were enforced.

Federated learning reduces the need to transmit raw customer, transaction, or network data. Soltani and colleagues use federated exchange of local IDS knowledge to address multi-sensor detection and traffic drift (Soltani et al., 2024). In heterogeneous cyber-physical systems, multi-agent federated reinforcement learning models data allocation and secure communication without requiring privacy-sensitive information to be shared, while addressing non-stationarity among devices (Xu et al., 2022). Privacy protection is not automatic: model updates can leak information, poisoned participants can corrupt shared knowledge, and institutional incentives affect participation. Distributed learning therefore requires secure aggregation, participant validation, provenance, and procedures for withdrawing compromised updates, issues emphasized in the broader trustworthy distributed-AI review (Wei & Liu, 2024).

Blockchain can support integrity and provenance, but it does not validate the truth of the data entered into the ledger. Research combining blockchain with explainable threat detection uses a permissioned proof-of-authority arrangement to validate and store data exchanged among cloud vendors before applying deep threat hunting (Kumar et al., 2024). The approach addresses auditability and source accountability, yet a validated record can still contain an erroneous or malicious observation. Cross-institutional intelligence should therefore preserve source identity, confidence, collection context, and dispute mechanisms. These requirements connect detection sharing to the policy and governance controls considered next.

3.5 Institutional Trust Boundaries in Federated Intelligence Exchange

Federated intelligence sharing requires a governance contract in addition to a communication protocol. Each participating institution should retain control of raw customer data, define approved update types, authenticate contributing agents, and apply minimum necessary disclosure. Shared artifacts should carry source and model version identifiers, creation timestamps, intended uses, integrity metadata, and withdrawal or revocation information.

Updates must be screened for poisoning, anomalous contribution patterns, and cross-domain incompatibility before aggregation. Participation agreements should allocate responsibility for erroneous intelligence, specify retention and deletion conditions, and establish procedures for dispute resolution and breach notification.

Such controls are necessary because secure transport and immutable logging protect the transmission and traceability of an artifact without guaranteeing its accuracy, lawfulness, or fitness for a specific financial decision.

4 DATA GOVERNANCE AND POLICY ENFORCEMENT ACROSS AUTONOMOUS FINANCIAL WORKFLOWS

4.1 Policy-aware agents for compliance automation, AML/KYC, and trade-finance controls

Autonomous security becomes operationally consequential when agents can change access, release transactions, modify data, or initiate investigations. **Policy-aware agents** constrain these actions by evaluating proposed operations against rules, authority levels, data classifications, and workflow context. In trade finance, the TradeQu architecture (Alaeifar et al., 2024) embeds a versioned and tenant-scoped Policy Store, a graph linking entities and documents, and policy gates within the multi-agent orchestrator. Immutable audit events record the resulting operations, allowing reviewers to reconstruct which rule governed a decision and which agent performed it.

This approach differs from post-hoc compliance review. AML and KYC workflows require continuous comparison of customer, transaction, document, and relationship evidence, while trade-finance controls may depend on document consistency, counterparty identity, and overlapping regulatory requirements. A knowledge graph can represent these relationships more directly than isolated document classifiers. The architecture should distinguish a recommendation, a screening result, a temporary hold, and an irreversible release. Each action has a different evidentiary threshold and should invoke a corresponding approval or escalation path.

The literature suggests that agentic finance can automate compliance and fraud tasks, but evidence remains uneven. A review of financial applications reports potential efficiency and compliance benefits alongside variable

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

evidence for real-world fraud detection and limited data on long-term cost effectiveness (Adeliyi, 2025). A proposed autonomous data-governance system combines performance optimization with encryption and PII obfuscation, reporting 99.9 percent adherence to SOC2 and GDPR protocols in its evaluation (Sanepalli, 2025). That result concerns the proposed framework and its testing conditions; it does not establish regulatory compliance across jurisdictions or institutional processes.

AML surveillance also illustrates why performance metrics must include operational cost. Financial-infrastructure research describes alert proliferation and false positives above 90 percent in expert AML domains (Alaeifar et al., 2024), motivating specialist agents and predictive maintenance. Reducing alerts without suppressing legitimate risk requires calibrated thresholds, case-level explanations, and feedback from investigators. Agents should learn from adjudicated cases only after data quality, label consistency, and potential bias have been assessed. A policy layer should block unauthorized model updates and preserve the rule version used at each decision.

4.2 Governance-as-a-service, zero-trust enforcement, and auditable policy execution

Governance-as-a-Service separates policy enforcement from the internal reasoning of individual agents. The proposed GaaS architecture intercepts agent outputs, evaluates declarative rules, assigns trust scores based on compliance and severity-weighted violations, and applies graduated interventions (Gaurav et al., 2025). Simulation regimes involving open-source models and financial decision-making indicate that the service can block or redirect high-risk behavior while maintaining throughput, although adversarial testing remains bounded by the selected models and scenarios. This separation permits common controls across heterogeneous agents and reduces reliance on voluntary agent cooperation.

Zero-trust architecture provides a complementary access principle: every request is evaluated according to current identity, context, and authorization rather than inherited network location. An agentic zero-trust proposal combines continuous authentication, behavioral analysis, threat hunting, policy enforcement points, human interfaces, and improvement loops (Gurram, 2025). In a financial workflow, the policy engine should evaluate the requesting agent, target resource, purpose, data sensitivity, transaction value, current threat state, and required human approvals. A successful prior action should not create indefinite authority.

Runtime governance also requires observable execution. The TradeQu design (Alaeifar et al., 2024) emits cryptographically verifiable audit events, while GaaS intercepts and logs actions for later analysis (Gaurav et al., 2025). These records should include the input evidence, model or policy version, retrieved sources, agent identity, delegated authority, proposed action, intervention outcome, and final approver. Such records support incident reconstruction and regulatory examination, but logging alone does not demonstrate that a decision was correct. Governance must connect records to testing, review, remediation, and the ability to suspend an agent.

A practical control model therefore combines preventive, detective, and corrective mechanisms. Preventive controls restrict tools, data, and action scope; detective controls monitor deviations, policy violations, and unusual delegation; corrective controls revoke credentials, quarantine agents, roll back changes, or require human review. The Model-Control-Policy framework similarly argues that autonomous workflows need explicit controls, explainability, human participation, and red-teaming (Suggu, 2025). This arrangement allows routine low-risk actions to proceed automatically while reserving high-impact decisions for supervised execution.

4.3 Formal Policy-Constrained Autonomous Decision Model

A practical distinction between adaptive recommendation and permissible execution can be expressed as a constrained action-selection problem. At decision time t, the coordinator receives an incident state, a collaborative risk estimate, and a set of candidate responses. The policy engine evaluates each proposed action against agent authority, asset criticality, data access rules, jurisdiction-specific obligations, service dependencies, and required approval status. The model selects a response only from the set of authorized options rather than attempting to compensate for a prohibited action with a favorable security reward.

The admissible action set is defined in Equation (5), which ensures that only responses satisfying the current institutional policy are considered for execution. The response coordinator then selects an optimal action according

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

to Equation (6), which minimizes a weighted combination of residual cyber exposure, service interruption, customer impact, and operational costs.

The decision is further constrained by Equation (7), which establishes a maximum acceptable probability of severe operational harm. If the admissible set is empty, or if potential harm cannot be estimated with acceptable confidence, the workflow defaults to escalation or an explicitly authorized safe action.

This formulation operationalizes the manuscript's central claim that regulatory and human authority are hard execution boundaries, not optional preferences within a learning objective. It also enables financial institutions to distinguish decision accuracy from execution permissibility, thereby reducing the possibility that a technically effective cybersecurity action produces unacceptable operational or regulatory consequences.

Equation (5): Admissible action set

$$\boxed{\mathcal{A}_t^{\text{allow}} = \{a \in \mathcal{A}_t : \text{Policy}(a, x_t, \mathcal{P}_t) = 1\}}$$

Where:

$\mathcal{A}_t^{\text{allow}}$ = set of authorized response actions.

$\mathcal{A}_t$ = set of candidate response actions.

a = an individual candidate action.

x_t = current incident state.

$\mathcal{P}_t$ = current institutional policy.

Policy$(\cdot)$ = authorization function, returning 1 for an authorized action.

Equation (6): Optimal response selection

$$a_t^* = \arg\min_{a \in \mathcal{A}_t^{\text{allow}}} [\alpha L_{\text{cyber}}(a, x_t) + \beta L_{\text{service}}(a, x_t) + \gamma L_{\text{customer}}(a, x_t) + \delta L_{\text{operation}}(a, x_t)]$$

Where:

a_t^* = selected response action.

L_{cyber} = residual cybersecurity loss.

L_{service} = service-disruption loss.

L_{customer} = customer-impact loss.

$L_{\text{operation}}$ = operational cost.

$\alpha, \beta, \gamma, \delta$ = institutional weighting parameters.

Equation (7): Operational safety constraint

$$\boxed{\Pr(\text{Harm}(a, x_t) > H_{\max}) \leq \varepsilon}$$

With the coefficient constraints:

$$\alpha, \beta, \gamma, \delta \geq 0$$
$$\alpha + \beta + \gamma + \delta = 1$$

Where:

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

$\text{Harm}(a, x_t)$ = estimated harm from action a

$H_{\max}$ = maximum acceptable harm threshold.

ε = maximum allowed probability of exceeding that threshold

4.4 Data privacy, provenance, explainability, and cross-organizational information sharing

Financial agents process identity records, payment histories, market information, corporate documents, and security telemetry. Data governance must therefore preserve purpose limitation, access control, retention rules, provenance, and separation between operational need and model-training need. Federated learning and synthetic data can reduce direct exposure: a privacy-oriented framework combines synthetic-data generation with differential privacy, encryption, and masking, reporting financial-sector accuracy of 94 percent and recall of 95 percent in its evaluation (Nadella et al., 2025). These results concern the framework's tested datasets and do not remove the need to assess re-identification, utility loss, and model leakage.

Provenance links an output to the data, rule, model, and agent that produced it. Blockchain-supported threat detection addresses integrity when cloud vendors exchange data, while explainable models provide analysts with reasons for a classification (Kumar et al., 2024). The two functions should not be conflated: an immutable record establishes what was submitted, whereas an explanation addresses why the system reached a decision. For customer screening or transaction holds, explanations should identify relevant evidence and policy conditions without revealing protected information or enabling evasion.

Cross-organizational sharing intensifies these constraints. Banks may need common intelligence about malware, compromised identities, or attack infrastructure, yet cannot disclose raw customer data or proprietary controls. Federated models, privacy-preserving updates, encrypted communication, and authenticated provenance offer technical options (Wei & Liu, 2024; Xu et al., 2022). Organizational governance must still specify who may contribute, inspect, challenge, or retract information. Data quality, liability, and accountability cannot be delegated to a communication protocol.

Responsible-AI research frames governance through structural, relational, and procedural practices, emphasizing the gap between general principles and their operational implementation (Papagiannidis et al., 2025). For distributed finance, this means defining ownership across model developers, security teams, compliance officers, infrastructure operators, and senior decision-makers. It also requires testing whether explanations are understood by investigators and whether privacy controls change detection performance across customer groups. These governance questions determine whether coordinated detection can safely feed the response mechanisms discussed below.

5 COORDINATED INCIDENT RESPONSE, AUTONOMOUS MITIGATION, AND SELF-HEALING OPERATIONS

5.1 Multi-agent response orchestration from detection through containment and recovery

Incident response becomes an orchestration problem when evidence arrives from many sensors and mitigation affects interdependent services. A detection agent may identify suspicious behavior, an intelligence agent may associate it with known techniques, an analysis agent may estimate scope, and a response agent may propose containment. MALCDF adopts this separation through four collaborating LLM agents and a secure communication layer. AutoBnB compares centralized, decentralized, and hybrid team structures in tabletop incident-response scenarios, indicating that configuration affects coordination and decision quality (Liu, 2025).

A governed response pipeline should move through evidence validation, risk assessment, action selection, authorization, execution, and recovery verification. Agents need shared incident identifiers, temporal state, confidence estimates, and explicit handoffs. A response controller should also distinguish containment from eradication and recovery, since disabling an account, isolating a workload, deleting persistence, and restoring service carry different risks. The response architecture must retain the evidence behind each transition so that later review can determine whether the action was proportionate.

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

Reinforcement learning can select sequential countermeasures when the attacker adapts. Hierarchical MARL decomposes defense decisions and reduces exploration difficulty in simulated network attack-defense games (Tang et al., 2024). Cloud defense research similarly groups defenders so they share information and jointly optimize resource allocation against changing attacker strategies (Chen et al., 2025). These approaches are suited to dynamic action selection, but their rewards must include service continuity, collateral effects, investigation cost, and policy compliance rather than only attack suppression.

Self-healing extends response beyond containment by restoring a secure operating state. RL-based self-healing proposals (Purves et al., 2024) combine changing threat perception, automated prescriptions, continuous learning, and federated multi-agent intelligence. Financial-infrastructure frameworks describe predictive analysis that identifies failure precursors and initiates restoration before customer-facing service loss. Such claims remain preliminary where evaluations use simulated enterprise networks or conceptual architectures. Recovery actions should therefore begin with reversible changes, tested rollback procedures, and explicit confirmation that security posture has improved.

Equation (8): Recovery validation function

$$\boxed{V_t = \mathbb{I}\left(R_{j,t+1} \leq \theta_R\right)\mathbb{I}(S_{t+1} \geq \theta_S)\mathbb{I}(C_{t+1} = 1)}$$

Where:

- V_t = recovery validation outcome.
- $\mathbb{I}(\cdot)$= indicator function returning 1 when the condition is satisfied.
- $R_{j,t+1}$ = post-action threat-risk score.
- θ_R = acceptable residual-risk threshold.
- S_{t+1} = post-action service health.
- θ_S = required service-health threshold.
- C_{t+1} = indicator that required compliance and integrity checks pass.

Successful recovery requires:

$$V_t = 1$$

If:

$$V_t = 0$$

the system must follow its approved failure-handling procedure rather than automatically declaring recovery.

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ **Research Article**

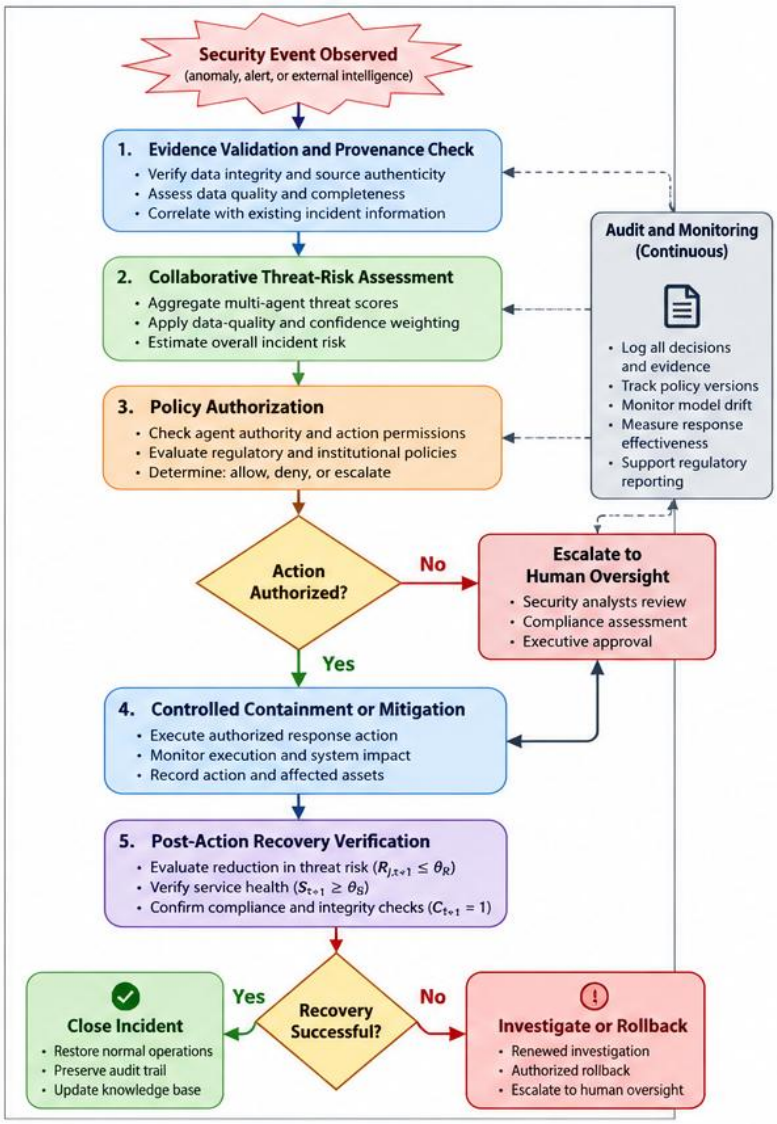


Figure 2. Policy-gated autonomous incident response lifecycle for distributed financial cybersecurity.

The workflow demonstrates the sequential stages of security event observation, evidence validation, collaborative risk assessment, policy authorization, controlled mitigation, and post-action recovery verification. Decision gates and human escalation mechanisms ensure that autonomous interventions remain compliant with institutional policies and operational safety requirements.

Figure 2 presents the proposed policy-gated autonomous incident response lifecycle, demonstrating how specialized AI agents coordinate security interventions while maintaining governance and operational accountability. The process begins when distributed observation agents identify a suspicious security event or anomalous operational condition. Collected evidence subsequently undergoes validation to assess its integrity, provenance, completeness, and reliability. Validated observations are aggregated through collaborative threat-risk assessment to estimate incident severity and determine the appropriate response requirements.

Before a mitigation action is executed, the policy authorization layer evaluates the proposed intervention against institutional security policies, agent permissions, operational constraints, and applicable governance requirements. Actions that satisfy the authorization criteria proceed to controlled containment or mitigation, while prohibited or high-impact actions are escalated for human review or rejected according to policy. Following execution, the recovery-verification mechanism evaluates residual threat exposure, service availability, and compliance or

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

integrity conditions. Successful verification permits incident closure and audit preservation, whereas unsuccessful verification initiates renewed investigation, authorized rollback, or human escalation.

Continuous audit monitoring records decisions, policy evaluations, agent activities, and operational outcomes throughout the lifecycle. This closed-loop mechanism establishes a distinction between threat identification, authorized response execution, and validated system recovery, thereby reducing the risks associated with unrestricted autonomous security operations in financially sensitive environments.

5.2 State-Based Incident Response and Recovery Verification

To prevent loosely coordinated agents from issuing contradictory or repeated commands, each security incident should follow a defined state-transition lifecycle. The proposed lifecycle comprises observation, evidence validation, risk assessment, policy authorization, controlled execution, recovery verification, and closure. Every transition records the incident identifier, initiating agent, evidence references, policy version, authorization status, action timestamp, and measured outcome. Concurrent actions that affect shared assets require idempotency controls and conflict resolution to avoid repeated containment or unintended service disruption.

An executed action is successful only when post-action observations demonstrate that the targeted risk has declined without unacceptable degradation in service or transaction integrity. A recovery-validation predicate can be expressed as Equation (8).

A failed validation triggers renewed investigation, an authorized rollback, or human escalation; it does not by itself authorize unrestricted retries. In payment and ledger environments, reversal must respect transaction finality and reconciliation rules: compensating transactions may be required where direct rollback would invalidate accounting or settlement records.

5.3 Cyber-risk mitigation in cloud, cyber-physical, and financial infrastructure environments

Risk-based mitigation is necessary because identical alerts can have different consequences across assets. A suspected compromise of a disposable test container may justify automatic isolation, while the same evidence near a payment ledger or trading system may require staged controls and human approval. The autonomous response controller for cyber-physical systems combines a hierarchical risk correlation tree with a competitive Markov decision process to estimate attacker paths and financial exposure (Kholidy, 2020). In a testbed involving an Aurora attack, it recovered the system after a cascading failure and reported faster mitigation than the compared Suricata response (Kholidy, 2020).

Cloud systems require agents to balance security against resource limits and service dependencies. GMADRLD models multiple attackers and defenders competing for cloud-server resources and uses grouped cooperation to manage computational complexity (Chen et al., 2025). A response agent that scales inspection or isolation must account for tenant boundaries, workload dependencies, and the possibility that containment will interrupt critical financial services. Policy-aware orchestration can encode those constraints before an RL policy selects among permitted actions.

Microservice environments create additional recovery problems because a transaction may depend on identity, authorization, messaging, database, and external API services. Organizational MARL research (Purves et al., 2024) includes an enterprise infrastructure and a microservices architecture among its application cases, using symbolic organizational constraints alongside learned policies. This combination is useful for defining which agents may restart, quarantine, or reconfigure a service. It also supports analysis of emergent roles after training, which is necessary when learned coordination differs from the designer's intended organization.

Financial risk mitigation must include economic and compliance consequences. An agent that blocks a suspicious payment can reduce fraud exposure but may delay legitimate settlement, breach service expectations, or create customer harm. A policy engine should therefore attach action costs to asset criticality, transaction context, uncertainty, and reversibility. Financial-system reviews identify promise in fraud detection and compliance automation but emphasize that empirical evidence remains limited and variable (Adeliyi, 2025). Response

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

evaluation should consequently report security outcomes alongside downtime, false interventions, recovery time, customer impact, and audit completeness.

5.4 Deception, adaptive defense, and self-healing responses against advanced persistent threats

Advanced persistent threats require defenders to manage incomplete knowledge and prolonged attacker adaptation. Cloud APT research models multiple attackers and defenders in an attack-defense game, then uses grouped multi-agent deep RL to share information within defender groups (Chen et al., 2025). This approach treats defense as an ongoing allocation problem rather than a single alert decision. Adaptive policies can alter monitoring intensity, isolate suspicious paths, or redirect attackers toward controlled environments, provided the action remains within legal and operational boundaries.

Cyber deception adds uncertainty to the attacker's decision process through honeypots, decoys, and honey tokens. Experiments using CyberBattleSim (Purves et al., 2024) report that deception reduced attacker cumulative rewards by up to 83.79 percent in a simulated ten-node topology, with effects varying by algorithm pairing. The study also acknowledges that reward penalties do not directly equal operational measures such as time to compromise or false-alert rates. Deception should therefore be evaluated through attacker dwell time, information gained, defender workload, and the risk that decoys expose sensitive architecture.

Self-healing systems combine adaptation with restoration. The response loop observes changes, updates a risk estimate, selects a permitted intervention, verifies system state, and feeds the result back into future decisions. Kholidy's controller demonstrates this logic through risk correlation, sequential game modeling, and recovery testing in a physical testbed (Kholidy, 2020). RL self-healing proposals (Purves et al., 2024) extend it with federated knowledge exchange, but the literature does not yet establish how continuously learned policies behave after a poisoning event or a major infrastructure change (Wei & Liu, 2024).

APT defense therefore needs a bounded autonomy model. Agents may automatically increase telemetry, deploy decoys, or quarantine low-criticality resources, while actions affecting settlement, market access, or physical safety require human authorization. Every intervention should be tested against rollback and recovery criteria, and learned policies should be frozen or replaced when their data, reward structure, or operating assumptions change. These safeguards connect adaptive response to the broader question of whether distributed autonomy remains secure and accountable.

6 TRUSTWORTHY AUTONOMY: SECURITY, ROBUSTNESS, EXPLAINABILITY, AND HUMAN ACCOUNTABILITY

6.1 Attack surfaces and vulnerabilities in LLM-based and agentic systems

Agentic systems expand the attack surface from model input to perception, memory, tools, communication, and execution. LLM reviews identify prompt injection, data poisoning, insecure output handling, adversarial instructions, and denial-of-service attacks (Ferrag et al., 2025). A lifecycle framework organizes related risks into perception failure, decision hallucination, memory poisoning, and malicious execution. In a financial SOC, a poisoned alert, manipulated retrieval result, or compromised tool response can redirect several cooperating agents, making the dependency chain more important than the accuracy of any single model.

Multi-agent coordination creates additional failure modes. Agents may accept false consensus, leak sensitive messages, amplify a mistaken assessment, or optimize a local reward at the expense of service continuity. Trustworthy distributed-AI research identifies evasion, poisoning, Byzantine behavior, irregular data distributions, privacy leakage, and fairness problems across distributed learning architectures (Wei & Liu, 2024). Secure channels reduce interception risk but do not ensure that an authenticated agent is honest or that its evidence is correct. Coordination protocols therefore need provenance, anomaly checks, disagreement handling, and the ability to isolate a malfunctioning participant.

6.2 Robustness, privacy, fairness, and secure coordination in distributed AI

Robustness must be evaluated during both training and deployment. Distributed systems can face poisoned updates, adversarial inference queries, and heterogeneous data that destabilizes convergence (Wei & Liu, 2024).

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

Federated reinforcement learning in heterogeneous cyber-physical systems addresses non-stationarity through multi-agent allocation and privacy-preserving interaction, but the need for incentive mechanisms shows that technical coordination depends on participant behavior (Xu et al., 2022). Financial institutions should test models against drift, missing data, malicious updates, and selective disclosure rather than rely on accuracy under benign benchmark conditions.

Privacy controls can reduce exposure while altering model utility. Differential privacy, encryption, masking, and synthetic data are proposed for enterprise security analytics, including financial applications (Nadella et al., 2025). Federated learning keeps raw observations local, but exchanged parameters or gradients may still reveal information and can carry poisoned behavior (Wei & Liu, 2024). Evaluation should report privacy budgets or attack resistance where applicable, as well as changes in recall, false positives, latency, and subgroup performance. A model that protects data but systematically misses risks for a customer segment is not trustworthy.

Fairness matters when agents influence identity verification, fraud screening, account restriction, or access decisions. The distributed-AI review treats fairness and governance as properties of both data and models (Wei & Liu, 2024). Financial institutions should examine whether proxy variables, uneven reporting, or historical enforcement patterns produce disparate error rates. Human review cannot correct this problem if explanations omit the relevant features or if investigators receive only a confidence score. Governance must therefore connect fairness testing to model updates, policy exceptions, and appeal procedures.

6.3 Explainable and human-supervised decision-making for high-stakes financial actions

Explainability has two audiences: technical operators who need to validate an alert and accountable decision-makers who must justify an intervention. Blockchain-enabled XAI research combines validated cloud data with attention-based threat hunting to support analyst decision-making (Kumar et al., 2024). Neuro-symbolic financial cyber-physical proposals (Li et al., 2024) similarly connect multi-agent detection with explainability, data governance, threat modeling, and feedback loops. These approaches suggest that explanations should identify evidence, causal or rule-based links, uncertainty, and the action alternatives considered.

Human supervision should be graduated rather than treated as a binary switch. Advisory agents can summarize evidence; supervised agents can execute pre-approved low-risk actions; autonomous agents can handle routine, reversible operations subject to monitoring. GaaS supports this arrangement through policy interception, trust scores, and graduated interventions (Gaurav et al., 2025). Zero-trust agentic architectures also retain human interfaces and oversight while automating continuous authentication and response orchestration (Gurram, 2025). The required level of supervision should increase with irreversibility, financial exposure, uncertainty, customer impact, and physical consequence.

Accountability must remain identifiable after delegation. Audit records should link each action to the responsible agent, policy version, evidence, authorization, and outcome. Responsible-AI governance research distinguishes structural arrangements, relationships among stakeholders, and procedures for implementing and evaluating principles (Papagiannidis et al., 2025). Applied to financial security, this means assigning responsibility for model development, data stewardship, policy approval, operational execution, and incident review. Governed autonomy is therefore not autonomy without people; it is delegated action within controls that preserve contestability, traceability, and the capacity to intervene.

7 CROSS-STUDY SYNTHESIS: COMPLEMENTARITIES, METHODOLOGICAL PATTERNS, AND CONTESTED DESIGN CHOICES

7.1 Evidence clusters spanning surveys, conceptual taxonomies, frameworks, and simulation-based evaluations

The literature converges on a layered model in which specialized agents detect, interpret, coordinate, enforce policy, and support response. Surveys and taxonomies establish the conceptual distinction between task-specific agents and agentic systems with decomposition, memory, tool use, and coordinated autonomy (Li et al., 2024), while cybersecurity reviews map these capabilities onto detection, response, and governance functions (Kaur et al., 2023). Primary studies then instantiate narrower components: continual distributed detection with federated

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

knowledge exchange (Soltani et al., 2024), hierarchical defense decisions in simulated networks (Tang et al., 2024), and runtime policy interception through Governance-as-a-Service (Gaurav et al., 2025).

Methodological patterns are equally clear. Evidence is concentrated in benchmark datasets, cyber ranges, tabletop exercises, testbeds, and simulations rather than operational financial environments. Reported improvements therefore establish technical feasibility under specified conditions, not reliable institution-wide outcomes. For example, Big-IDS reports 97.44% accuracy on NSL-KDD (Louati et al., 2024), whereas MALCDF reports 90.0% accuracy, 85.7% F1-score, 9.1% false-positive rate, and 6.8-second event latency on a 50-record stream derived from CICIDS2017. These results are not directly comparable because their datasets, baselines, task definitions, and evaluation scales differ.

7.2 Distributed autonomy versus centralized oversight and control

Distributed autonomy addresses the scale, heterogeneity, and partial observability of financial infrastructure. Local agents can process data near hosts, services, or institutional boundaries, while coordination supports collective detection and reduces dependence on a single control point. Federated learning further permits knowledge exchange without transferring raw data, although heterogeneous data and nonstationary environments complicate convergence and trust (Wei & Liu, 2024; Xu et al., 2022).

Centralized oversight remains necessary where actions affect customers, liquidity, access, or regulatory obligations. Hierarchical MARL reduces action-space complexity by separating strategic coordination from local defense behavior (Tang et al., 2024), while hybrid organizational models constrain learned policies through explicit roles and obligations. The defensible design is therefore federated execution with centralized authority over policy, escalation, evidence retention, and exceptional actions, rather than either complete local independence or a single centralized decision-maker.

7.3 Performance gains versus auditability, interpretability, and operational risk

The strongest performance claims concern speed, adaptability, and false-positive reduction. Multi-agent threat-intelligence research reports higher detection accuracy, lower response latency, and fewer false positives than its centralized comparison, while enterprise threat-hunting work reports rapid event processing and automated SOAR workflows (Swarna, 2025). Yet these outcomes do not establish safety: faster autonomous action can amplify an erroneous classification, corrupt shared knowledge, or interrupt legitimate financial activity.

Auditability must consequently be designed into communication, memory, and execution. Immutable events, policy versions, agent identities, and evidence references support retrospective review in trade-finance orchestration, and explainable threat hunting can assist analysts when evidence from multiple cloud sources is validated (Kumar et al., 2024). Runtime governance adds interception, trust scoring, graduated intervention, and logging without requiring changes to model internals (Gaurav et al., 2025). The central design choice is thus not autonomy versus no autonomy, but which actions are reversible, which require approval, and what evidence must accompany each decision.

7.4 Proposed Evaluation Protocol and Comparative Baselines

The integrated reference architecture should be evaluated through an explicitly staged experimental design. The first stage uses replayable network and transaction-event streams with known ground truth to establish detection behavior. The second stage introduces controlled adversarial scenarios, delayed and missing telemetry, compromised agents, poisoned updates, and conflicting observations. The third stage tests response execution in a non-production financial service environment or cyber range with realistic transaction dependencies, rollback procedures, and synthetic or appropriately de-identified records. Human oversight and policy enforcement must be enabled throughout action-execution experiments. Production experiments, if later authorized, require separate institutional review and documented operational safeguards.

Four comparison conditions are recommended: (i) centralized rule-based detection and human-led response; (ii) centralized machine-learning detection with conventional response orchestration; (iii) distributed multi-agent detection without a separate governance authorization layer; and (iv) the proposed governance-aware multi-agent

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

approach. Ablation experiments should independently disable provenance weighting, quality weighting, cross-agent coordination, and recovery verification, while retaining essential safety controls. All baselines must process comparable workloads and be evaluated under identical incident definitions, cost assumptions, and hardware conditions.

The evaluation should report precision, recall, F1-score, area under the precision–recall curve where appropriate, false-positive alerts per workload, detection delay, policy-violation attempts blocked, unauthorized actions executed, mean time to contain, mean time to recover, transaction disruption, customer-impact indicators, audit-record completeness, and analyst intervention frequency.

Operational governance effectiveness should be measured separately from prediction quality. Results should include confidence intervals or repeated-trial variability and disclose dataset leakage checks, class prevalence, threshold calibration, and the cost of false interventions. These measures are a proposed protocol and must not be presented as experimental findings until implementation and testing are completed.

Table 1. Representative multi-agent cybersecurity studies: settings, findings, and limitations.

Study	Application setting	Agent architecture	Evaluation environment	Reported findings	Governance controls	Limitations
Louati et al. (2024)	Distributed intrusion detection	Decentralized multi-agent RL	NSL-KDD	97.44% accuracy	Not reported	Dataset generalizability
Tang et al. (2024)	Network defense	Hierarchical MARL	Simulated cyber defense	Coordinated defense strategies	Hierarchical role separation	Simulation-to-production gap
Gaurav et al. (2025)	Runtime AI governance	Governance-as-a-Service	Simulation	Policy interception and graduated interventions	Explicit policy enforcement	Bounded evaluation scenarios

Equation (9): Detection performance metrics

Precision:

$$\text{Precision} = \frac{TP}{TP + FP}$$

Recall:

$$\text{Recall} = \frac{TP}{TP + FN}$$

F1-score:

$$\boxed{F_1 = \frac{2TP}{2TP + FP + FN}}$$

Where:

- TP= true positive detections.
- FP = false positive detections.
- FN = false negative detections.

Metrics with zero denominators require an explicitly documented reporting convention.

Equation (10): Audit completeness

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

$$A_{\text{complete}} = \frac{N_{\text{actions with all required audit fields}}}{N_{\text{executed actions}}}$$

Where:

- A_{complete} = audit record completeness rate.
- Numerator = executed actions having all required audit fields.
- Denominator = total number of executed actions.

If no actions were executed, the completeness ratio is undefined and should be reported accordingly rather than automatically assigned a perfect score.

Table 2. Recommended evaluation dimensions and metrics.

Dimension	Recommended metrics
Threat detection	Precision, recall, F1-score, PR-AUC
Response efficiency	Detection delay, containment time, recovery time
Operational reliability	Service availability, transaction disruption
Governance	Unauthorized actions, policy violation blocking
Auditability	Audit completeness, evidence traceability
Privacy and fairness	Information leakage, subgroup error rates
Resource efficiency	CPU, memory, communication overhead
Human oversight	Escalation frequency, override performance

Table 3. Recommended experimental comparison.

Experimental condition	Detection architecture	Governance mechanism
Baseline A	Centralized rule-based detection	Human-led response
Baseline B	Centralized machine learning	Conventional response orchestration
Baseline C	Distributed multi-agent detection	No separate proposed policy authorization layer
Proposed framework	Distributed multi-agent detection	Explicit policy authorization, oversight, and recovery verification

These are recommended experimental groups, not results obtained in the current paper. All live action execution must remain subject to essential safety controls, including in comparative tests.

8 RESEARCH GAPS FOR AN INTEGRATED MULTI-AGENT FINANCIAL CYBERSECURITY FRAMEWORK

8.1 Limited empirical validation in live, heterogeneous, and regulated financial environments

Current evidence rarely tests the full operational conditions of finance: heterogeneous legacy systems, cloud services, market workloads, privacy constraints, institutional boundaries, and supervisory scrutiny. Financial reviews describe promising applications in compliance, fraud detection, and smart contracts, but report limited

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

real-world and long-duration evidence (Adeliyi, 2025). Benchmark studies also depend heavily on NSL-KDD, CICIDS2017, CICIDS2018, or simulated environments (Louati et al., 2024; Soltani et al., 2024; Tang et al., 2024). Future evaluation requires prospective pilots, heterogeneous datasets, drift monitoring, recovery outcomes, analyst workload, customer impact, and regulatory audit results.

8.2 Missing integration between threat detection, policy enforcement, and incident response

Research commonly evaluates detection, governance, or response as separate capabilities. Detection agents generate scores, policy systems constrain outputs, and response agents execute containment, but few studies demonstrate a shared state, common evidence model, and closed feedback loop across all three functions. Framework proposals identify this integration as a gap in financial infrastructure (Sanepalli, 2025), while agentic cybersecurity reviews warn that tool-use correctness, long-horizon reasoning, and response validation remain unresolved. An integrated framework should therefore test whether detection confidence, policy constraints, and response consequences remain consistent during evolving incidents.

8.3 Unresolved standards for accountability, interoperability, adversarial testing, and governance evaluation

The field lacks agreed measures for attributing a distributed decision, comparing agent protocols, and evaluating governance beyond task accuracy. Robustness research identifies poisoning, Byzantine behavior, privacy, fairness, and inference attacks as distinct requirements (Wei & Liu, 2024), while LLM reviews document prompt injection, insecure output handling, and adversarial instructions (Ferrag et al., 2025). Governance research further separates structural, relational, and procedural responsibilities (Papagiannidis et al., 2025). Standards should connect these dimensions to interoperable message schemas, red-team protocols, escalation records, explanation quality, and measurable human override performance.

8.4 Contributions, Boundaries, and Limitations of the Proposed Framework

This review contributes an integrated analytical blueprint in four respects. First, it connects distributed detection, multi-agent coordination, policy enforcement, and incident recovery within one reference architecture. Second, it proposes a common evidence representation that makes data quality, source confidence, and provenance explicit inputs to collaborative risk assessment. Third, it separates adaptive action recommendation from authorized execution through a policy-constrained decision model and verifiable recovery criteria. Fourth, it defines an evaluation protocol that measures governance compliance and operational consequences alongside predictive accuracy.

These contributions are conceptual and methodological; they do not establish a novel trained algorithm or demonstrate superior performance in live financial institutions. Limitations include heterogeneous evidence sources, reliance on simulation and benchmark reports, the absence of a common financial incident dataset, and unresolved questions about cross-institutional interoperability, economic cost, fairness, and regulatory variation.

Accordingly, the proposed architecture should be treated as a set of testable design propositions whose effectiveness and feasibility require implementation, comparative experiments, adversarial validation, and longitudinal deployment studies.

9 CONCLUSION: DESIGNING GOVERNABLE MULTI-AGENT SECURITY FOR DISTRIBUTED FINANCE

9.1 Synthesis of the evidence-based framework

The evidence supports a **governable multi-agent architecture** with local sensing, federated intelligence, hierarchical coordination, policy gates, human escalation, and auditable execution. Detection agents should combine behavioral, symbolic, and language-based evidence; policy agents should constrain actions at runtime; response agents should prefer reversible containment before irreversible intervention. Secure communication, provenance, explainability, and continuous evaluation must extend across the complete perception-to-action chain (Gaurav et al., 2025).

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ **Research Article**

9.2 Implications for research, financial institutions, and regulators

Researchers should move from isolated benchmark claims toward comparative, adversarial, and longitudinal evaluation in realistic financial settings. Institutions should begin with bounded workflows, graduated autonomy, immutable records, and explicit accountability for models, policies, data, and operators. Regulators should assess evidence quality, override capacity, cross-institutional interoperability, and incident reporting, rather than treating autonomy as a sufficient indicator of security. Governed delegation offers the practical basis for distributed financial defense because it preserves local responsiveness while retaining institutional control. Future work should implement the proposed reference architecture in a controlled financial cyber range and evaluate the full evidence-to-action chain under benign, adversarial, and degraded-data conditions. Particular attention should be given to calibrated cross-agent risk aggregation, policy conflicts across jurisdictions, failure containment after an agent compromise, financial transaction consistency during rollback, and the cost of human escalation.

Comparative testing should determine not only whether coordination improves detection and response but also whether the governance layer measurably reduces unsafe interventions without unacceptable latency or service disruption. This empirical agenda would establish the conditions under which governed multi-agent autonomy offers a defensible operational advantage over conventional centralized security orchestration.

DECLARATIONS

Author contributions: David Excel Ozowara: Conceptualization, Writing – original draft, Writing – review & editing. Chukwudera Obumneke Anunagba: Conceptualization, Writing – original draft, Writing – review & editing.

Funding: This research received no external funding.

Competing interests: The authors declare no competing interests.

Data availability: No new data were generated or analysed in this study.

Use of generative AI: Generative AI tools were used to assist with drafting and language editing. The authors reviewed and edited all content and take full responsibility for the manuscript.

REFERENCES

[1] Adeliyi, O. A. (2025). Agentic AI in financial systems: Architectures, compliance automation, fraud detection, and smart contract orchestration. *Figshare*. https://doi.org/10.6084/m9.figshare.29891762

[2] Alaeifar, P., Pal, S., Jadidi, Z., Hussain, M., & Foo, E. (2024). Current approaches and future directions for cyber threat intelligence sharing: A survey. Journal of Information Security and Applications, 83, Article 103786. https://doi.org/10.1016/j.jisa.2024.103786

[3] Chen, J., Lan, X., Zhang, Q., Ma, W., Fang, W., & He, J. (2025). Defending against APT attacks in cloud computing environments using grouped multiagent deep reinforcement learning. *IEEE Internet of Things Journal*, *12*(12), 19459–19470. https://doi.org/10.1109/jiot.2025.3542119

[4] Ferrag, M. A., Alwahedi, F., Battah, A. A., Cherif, B., Mechri, A., Tihanyi, N., Bisztray, T., & Debbah, M. A. (2025). Generative AI in cybersecurity: A comprehensive review of LLM applications and vulnerabilities. *Internet of Things and Cyber-Physical Systems*, *5*, 1–46. https://doi.org/10.1016/j.iotcps.2025.01.001

[5] Finistrella, S., Mariani, S., & Zambonelli, F. (2025). Multi-Agent reinforcement learning for cybersecurity: Classification and survey. *Intelligent Systems with Applications*, *26*, 200495. https://doi.org/10.1016/j.iswa.2025.200495

[6] Gaurav, S., Heikkonen, J., & Chaudhary, J. (2025). Governance-as-a-Service: A multi-agent framework for AI system compliance and policy enforcement [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2508.18765

[7] Gurram, A. (2025). Generative AI for enhanced cybersecurity: building a zero-trust architecture with agentic AI. *World Journal of Advanced Engineering Technology and Sciences*, *15*(1), 2380–2396. https://doi.org/10.30574/wjaets.2025.15.1.0504

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

[8] He, M., Wang, X., Wei, P., Liu, Y., Teng, Y., & Lyu, R. (2024). Reinforcement learning meets network intrusion detection: A transferable and adaptable framework for anomaly behavior identification. *IEEE Transactions on Network and Service Management*, *21*(2), 2477–2492. https://doi.org/10.1109/tnsm.2024.3352586

[9] Kaur, R., Gabrijelčič, D., & Klobučar, T. (2023). Artificial intelligence for cybersecurity: Literature review and future research directions. *Information Fusion*, *97*, 101804. https://doi.org/10.1016/j.inffus.2023.101804

[10] Kholidy, H. A. (2020). Autonomous mitigation of cyber risks in the cyber–physical systems. *Future Generation Computer Systems*, *115*, 171–187. https://doi.org/10.1016/j.future.2020.09.002

[11] Kiely, M., Ahiskali, M., Borde, É., Bowman, B., Bowman, D. B., Van Bruggen, D., Cowan, K., Dasgupta, P., Devendorf, E., Edwards, B., Fitts, A., Fugate, S. J., Gabrys, R., Gould, W., Huang, H., Jacobs, J., Kerr, R., King, I. J., Li, L., … Zhang, Y. (2025). Exploring the efficacy of multi-agent reinforcement learning for autonomous cyber defence: A CAGE challenge 4 perspective. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 39, Issue 28, pp. 28907–28913). Association for the Advancement of Artificial Intelligence (AAAI). https://doi.org/10.1609/aaai.v39i28.35158

[12] Kumar, P., Javeed, D., Kumar, R., & Islam, A. K. M. N. (2024). Blockchain and explainable AI for enhanced decision making in cyber threat detection. *Software Practice and Experience*, *54*(8), 1337–1360. https://doi.org/10.1002/spe.3319

[13] Landolt, C. R., Würsch, C., Meier, R., Mermoud, A., & Jang-Jaccard, J. (2025). Multi-Agent reinforcement learning in cybersecurity: From fundamentals to applications [Preprint]. In *arXiv (Cornell University)*. https://doi.org/10.48550/arxiv.2505.19837

[14] Li, X., Wang, S., Zeng, S., Wu, Y., & Yang, Y. (2024). A survey on LLM-based multi-agent systems: Workflow, infrastructure, and challenges. Vicinagearth, 1, Article 9. https://doi.org/10.1007/s44336-024-00009-2

[15] Li, Y., Fang, H., & Chen, J. (2020). Anomaly detection and identification for multiagent systems subjected to physical faults and cyberattacks. *IEEE Transactions on Industrial Electronics*, *67*(11), 9724–9733. https://doi.org/10.1109/tie.2019.2952802

[16] Liu, Z. (2025). AutoBnB: Multi-Agent incident response with large language models. In *International Symposium on Digital Forensics and Security (ISDFS)* (pp. 1–6). https://doi.org/10.1109/isdfs65363.2025.11012055

[17] Louati, F., Ktata, F. B., & Amous, I. (2024). Big-IDS: a decentralized multi agent reinforcement learning approach for distributed intrusion detection in big data networks. *Cluster Computing*, *27*(5), 6823–6841. https://doi.org/10.1007/s10586-024-04306-9

[18] Nadella, G. S., Addula, S. R., Yadulla, A. R., Sajja, G. S., Meesala, M. K., Maturi, M. H., Meduri, K., & Gonaygunta, H. (2025). Generative AI-Enhanced cybersecurity framework for enterprise data privacy management. *Computers*, *14*(2), 55. https://doi.org/10.3390/computers14020055

[19] Olowu, O., Adeleye, A. O., Omokanye, A. O., Ajayi, A. M., Adepoju, A. O., Omole, O. M., & Chianumba, E. C. (2024). AI-driven fraud detection in banking: A systematic review of data science approaches to enhancing cybersecurity. *GSC Advanced Research and Reviews*, *21*(2), 227–237. https://doi.org/10.30574/gscarr.2024.21.2.0418

[20] Ouiazzane, S., Addou, M., & Barramou, F. (2021). A multiagent and machine learning based hybrid NIDS for known and unknown cyber-attacks. *International Journal of Advanced Computer Science and Applications*, *12*(8). https://doi.org/10.14569/ijacsa.2021.0120843

[21] Papagiannidis, E., Mikalef, P., & Conboy, K. (2025). Responsible artificial intelligence governance: A review and research framework. *The Journal of Strategic Information Systems*, *34*(2), 101885. https://doi.org/10.1016/j.jsis.2024.101885

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

[22] Park, T. (2024). Enhancing anomaly detection in financial markets with an LLM-based multi-agent framework [Preprint]. In *RePEc: Research Papers in Economics*. https://doi.org/10.48550/arxiv.2403.19735

[23] Purves, T., Kyriakopoulos, K. G., Jenkins, S., Phillips, I., & Dudman, T. (2024). Causally aware reinforcement learning agents for autonomous cyber defence. Knowledge-Based Systems, 304, 112521. https://doi.org/10.1016/j.knosys.2024.112521

[24] Sanepalli, U. R. (2025). Autonomous medallion orchestration: A multi-agent reinforcement learning framework for financial ecosystems. *International Journal For Multidisciplinary Research*, *7*(2). https://doi.org/10.36948/ijfmr.2025.v07i02.69007

[25] Sethi, K., Madhav, Y. V., Kumar, R., & Bera, P. (2021). Attention based multi-agent intrusion detection systems using reinforcement learning. *Journal of Information Security and Applications*, *61*, 102923. https://doi.org/10.1016/j.jisa.2021.102923

[26] Shi, G., & He, G. (2021). Collaborative multi-agent reinforcement learning for intrusion detection. In *IEEE International Conference on Network Intelligence and Digital Content (IC-NIDC)* (pp. 245–249). https://doi.org/10.1109/ic-nidc54101.2021.9660402

[27] Soltani, M., Khajavi, K., Siavoshani, M. J., & Jahangir, A. H. (2024). A multi-agent adaptive deep learning framework for online intrusion detection. *Cybersecurity*, *7*(1). https://doi.org/10.1186/s42400-023-00199-0

[28] Suggu, S. K. (2025). Agentic AI workflows in cybersecurity: Opportunities, challenges, and governance via the MCP model. *Journal of Information Systems Engineering & Management*, *10*(52s), 612–624. https://doi.org/10.52783/jisem.v10i52s.10767

[29] Swarna, R. (2025). *Enterprise AI-Enabled multi-agent threat hunting: A comprehensive framework for autonomous cybersecurity operations* [Preprint]. Institute of Electrical and Electronics Engineers (IEEE). https://doi.org/10.36227/techrxiv.175825523.37911598/v1

[30] Tang, Y., Sun, J., Wang, H., Deng, J., Liang, T., & Xu, W. (2024). A method of network attack-defense game and collaborative defense decision-making based on hierarchical multi-agent reinforcement learning. *Computers & Security*, *142*, 103871. https://doi.org/10.1016/j.cose.2024.103871

[31] Wei, W., & Liu, L. (2024). Trustworthy distributed AI systems: Robustness, privacy, and governance. *ACM Computing Surveys*, *57*(6), 1–42. https://doi.org/10.1145/3645102

[32] Xu, M., Peng, J., Gupta, B. B., Kang, J., Xiong, Z., Li, Z., & El-Latif, A. A. A. (2022). Multiagent federated reinforcement learning for secure incentive mechanism in intelligent cyber–physical systems. *IEEE Internet of Things Journal*, *9*(22), 22095–22108. https://doi.org/10.1109/jiot.2021.3081626

Copyright © 2025 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.