INTERNATIONAL JOURNAL OF ACTA INFORMATICA
VOLUME (2026)

IDENTITY-CONSCIOUS GOVERNANCE FOR AUTONOMOUS AI AGENTS: A FRAMEWORK FOR ENTERPRISE AUTHORIZATION, DELEGATION, AND AUDITING

Muhammad Salah Ud Din[1]

[1]Jagnnath Univeristy, BANGLADESH

ABSTRACT

AI agent systems capable of autonomous operation in enterprise environments are swiftly evolving from research prototypes to production implementations. These technologies acquire ambient credentials, function non-deterministically, and lack standardised methods for associating actions with human identity, resulting in a security stance that is fundamentally incompatible with organisational governance standards. This paper introduces a thorough framework for identity-aware agent governance that addresses five significant deficiencies in existing architectures: (1) the eradication of ambient authority via identity-inherited permissions, (2) infrastructure-level policy enforcement that is independent of model compliance, (3) sequence-aware authorisation capable of identifying composition attacks within the authorised scope, (4) independent action verification that tackles the oracle problem of agents self-reporting their actions, and (5) hallucination-aware audit systems that consider the unreliability of LLM reasoning. Our analysis is based on documented vulnerabilities from production agent systems, including the ClawHavoc supply chain attack (over 824 malicious agent skills), Cisco's discovery that 26% of analysed agent skills exhibit security vulnerabilities, and an architectural examination of credential inheritance patterns in Claude Code Agent Teams, Cursor FastRender (2,000 concurrent agents), and OpenClaw (over 30,000 internet-exposed instances). We assess the relevance of current standards (OAuth 2.0, SPIFFE/SPIRE, Zero Trust Architecture, MCP) and pinpoint particular changes necessary for non-deterministic agent processes. Our primary contribution is a stratified identity model utilising cryptographic delegation tokens that disseminate human authorisation throughout multi-agent chains, facilitating per-call policy assessment at the infrastructure level.

KEYWORDS: Identity-Conscious Governance; Autonomous AI Agents; Enterprise Authorization; Delegation; Auditing

INTRODUCTION

The swift implementation of AI agent systems capable of autonomous operation signifies a significant transformation in industrial computing. In contrast to conventional software that follows predetermined code routes, AI agents autonomously

determine the tools to utilise, their sequence, and the settings to apply, based on reasoning that fluctuates with each input. The non-determinism, coupled with the lack of identity-aware authorisation procedures, results in a security posture that is incompatible with the governance requirements of organisations in regulated industries.

Contemporary agent frameworks — such as GitHub Copilot, Claude Code, Cursor, Google Antigravity, and OpenClaw — exhibit a uniform architectural paradigm: the agent functions with the credentials available in the runtime environment, irrespective of the individual who initiated the activity or their authorisation level. The concept we refer to as ambient authority inheritance indicates that the prompts of a junior engineer and a leading architect operate under the same permissions. The notion of ambient authority is firmly rooted in capability-based security literature [Dennis and Van Horn, 1966; Miller et al., 2003], although its implementation in LLM-based agent systems presents new dangers due to the agent's action space being variable post-deployment.

THIS MANUSCRIPT PRESENTS THE SUBSEQUENT CONTRIBUTIONS:

A layered identification framework for AI agents that distinguishes between agent type identity, instance identity, delegated human identity, and session/task identity, incorporating cryptographic binding across the layers.

An identity-aware authorisation framework that does per-call policy evaluation at the infrastructure layer, irrespective of model compliance, enhanced with sequence-aware analysis to identify composition attacks.

• Examination of the oracle problem in agent systems, wherein the agent concurrently serves as both actor and observer of its own behaviours, along with suggested independent verification methodologies.

• Recognition of hallucination as a primary security concern in agentic systems, separate from adversarial attacks, affecting the trustworthiness of audit trails.

• Empirical foundation in recorded vulnerabilities from production agent systems, encompassing supply chain assaults, credential disclosure, and architectural examination of multi-agent credential inheritance.

CONTEXT AND RELEVANT LITERATURE

ARCHITECTURES OF AI AGENTS

AI agent systems comprise a minimum of one generative AI model and supporting software that provides the model with tools to execute various discretionary actions [NIST, 2026]. The standard architecture is as follows: user initiation → agent orchestrator → reasoning model (LLM) → tool invocations → external systems. Multi-agent architectures enhance this through agent-to-agent communication, hierarchical job decomposition, and swarm coordination patterns.

The Model Context Protocol (MCP) [Anthropic, 2024] has arisen as a standardisation framework for agent-to-tool communication, facilitating tool discovery and invocation methodologies. Currently, MCP standardises the available tools, but not the authorisation of their usage or the extent of such usage. The protocol has commenced the integration of OAuth 2.1 for authorisation; however, its implementation is still nascent and uneven among frameworks.

CAPABILITY-ORIENTED SECURITY AND AMBIENT AUTHORITY

The principle of least authority (POLA) and the risks associated with ambient authority are well-documented in capability-based security literature [Dennis and Van Horn, 1966; Miller et al., 2003]. In capability systems, access permissions are expressly assigned rather than derived from the surrounding context. Contemporary AI agent architectures inherently contravene the Principle of Least Authority (POLA): agents acquire all credentials present in their runtime environment, resulting in what practitioners refer to as "god-mode API keys," akin to the academic concept of ambient authority [Stiegler, 2004].

SECURITY OF LARGE LANGUAGE MODELS

Adversarial assaults on language models have been thoroughly recorded. Indirect quick injection [Greshake et al., 2023] illustrates that data handled by the model can supersede system-level directives. Universal adversarial prefixes [Zou et al., 2023] offer automated techniques for circumventing instruction-following behaviour. The OWASP Top 10 for LLM Applications [OWASP, 2025] enumerates model-layer vulnerabilities such as prompt injection, unsafe output management, and training data contamination.

Research on misleading alignment [Hubinger et al., 2024] has shown that fine-tuned models can incorporate enduring backdoor behaviours ("sleeper agents") that activate upon specified triggers while successfully passing conventional safety assessments. This discovery has immediate ramifications for agent systems utilising fine-tuned models, as training-time assaults pose a supply chain risk at the model level.

ZERO TRUST ARCHITECTURE

NIST SP 800-207 [Rose et al., 2020] delineates the zero trust principle: never trust, always verify. The fundamental notion of per-request authorisation based on numerous signals, although intended for network topologies, directly applies to agent systems. The standard presupposes deterministic agents with predictable access patterns, necessitating an extension for non-deterministic agents that dynamically identify and utilise tools.

AGENT IDENTITY STANDARDS

The NIST NCCoE concept paper on Software and AI Agent Identity and Authorisation

[NIST NCCoE, 2026] delineates the primary challenge: current identity standards (OAuth 2.0, OpenID Connect, SPIFFE/SPIRE, SCIM) were formulated for human users and deterministic software services. Agents necessitate novel frameworks for articulating delegation, transmitting human identity across multi-step procedures, and linking autonomous activities to human authorisation.

THREAT MODEL

We delineate seven categories of security threats that are either unique to or exacerbated in AI agent systems, based on verified vulnerabilities from production deployments.

AMBIENT AUTHORITY AND CREDENTIAL INHERITANCE

In Claude Code Agent Teams, all subordinate agents adopt the authorisation settings of the lead agent, encompassing the

--dangerously-skip-permissions flag. In Cursor's FastRender architecture, 2,000 concurrent agents function with distinct role delineation (architects, managers, workers, judges) while maintaining a unified scope of credentials. A worker agent possesses the same system access as an architect agent. Role separation without permission separation facilitates coordination but neglects security. More than 30,000 OpenClaw instances have been detected as publicly accessible, while the Moltbook database has revealed 1.5 million API tokens and 35,000 email addresses.

PROMPT INJECTION AS A MEANS OF CIRCUMVENTING AUTHORISATION

When authorisation policies are articulated as natural language directives within the model's context window, prompt injection not only compromises output quality but also completely circumvents authorisation. Retrieval-augmented generation (RAG) systems exacerbate this risk: papers retrieved during a job are integrated into the same context window as system-level policies, and a hacked document inside a RAG corpus can introduce instructions that are indistinguishable from authentic directives. In corporate settings, accessing several data sources renders each retrieval a possible injection vector.

COMPOSITION ATTACKS WITHIN PERMITTED PARAMETERS

An agent can execute multi-step actions in which each step undergoes per-call authorisation checks, yet the overall sequence results in an unauthorised output. A file read and a network call, both individually authorised, provide data exfiltration within the agent's authorised parameters. In contrast to conventional software, where execution routes are predictable and can be audited beforehand, an agent's emergent reasoning might generate innovative multi-step tactics that are not foreseen by any established policy rule. This signifies a fundamental constraint of per-call authorisation devoid of sequence-aware analysis.

HALLUCINATION AS A SECURITY THREAT

Unlike adversarial attacks, hallucination presents immediate security threats: agents fabricate tool parameters (invoking genuine APIs with erroneous inputs like incorrect account numbers or file paths that lead to unintended resources), misinterpret authorisation (concluding they possess permissions they lack), and generate chain-of-thought explanations that serve as post-hoc rationalisations rather than true representations of their reasoning process. This final attribute compromises the credibility of the audit trail, as documenting the agent's self-reported rationale does not ensure an accurate account of the reasons for activities made.

THE ORACLE DILEMMA

The agent concurrently serves as both actor and observer of its own actions. An agent may report success despite having failed silently, hallucinated the action, or executed an unintended action. Current ideas for audit and non-repudiation necessitate that the agent generates a record of its reasoning chain, without independent verification to ensure that the described action aligns with the accomplished action. This establishes a critical deficiency in accountability frameworks.

SUPPLY CHAIN ASSAULTS

The ClawHavoc campaign, recognised by Koi Security, recorded over 824 malicious skills in the OpenClaw ClawHub marketplace, encompassing infostealers, keyloggers, and backdoors, with one supply chain attack responsible for 335 infections (CVE-2026-25253). Cisco's investigation revealed that 26% of 31,000 assessed agent skills exhibited vulnerabilities, including 2 critical and 5 high-severity vulnerabilities inside a single skill that facilitated quiet data exfiltration. Supply chain hazards permeate the model layer: meticulously calibrated models may incorporate enduring backdoor behaviours [Hubinger et al., 2024], and reinforcement learning from human feedback (RLHF) training can instill systematic biases in adherence to authorisation instructions.

MULTI-AGENT AMPLIFICATION

Multi-agent systems exacerbate existing threats via credential scope inheritance across agent boundaries, delegation chain opacity (resulting in the loss of the original user's identity during agent-to-agent transitions), cross-context leakage (where sensitive data disseminates through agent summaries), inter-agent prompt injection (facilitating lateral movement within swarms), data aggregation risk (where agents collectively compile information surpassing any individual user's access level), and recursive cost amplification (where adversarial inputs generate recursive sub-agents).

PROPOSED FRAMEWORK: GOVERNANCE OF IDENTITY-AWARE AGENTS

LAYERED IDENTITY FRAMEWORK

We propose a four-tier identification framework for corporate agent implementations:

• Layer 1 — Agent Type Identity: Classification of the agent (coding assistance, data analyst, workflow executor). Established at deployment, recorded in the enterprise agent registry. Comparable to service catalogue entries.

• Layer 2 — Agent Instance Identity: Cryptographic identity for each active instance, consistent with SPIFFE/SPIRE workload identity. Comprises agent type/version, deploying organization, runtime environment, and capability manifest.

• Layer 3 — Delegated Human Identity: The authority under which the agent functions. Each agent activity possesses a provable relationship to the initiating human identity. This constitutes the essential absent component in all existing agent architectures.

• Layer 4 — Session/Task Identity: Temporary identity confined to a certain task or dialogue. Facilitates permission scoping to specific aims instead of universal access.

Layers 1 and 2 are permanent; Layers 3 and 4 are transient. The transient layers are cryptographically linked to the static layers, establishing a verifiable sequence: agent instance → authorised by user → for scope → at time.

IDENTITY-AWARE AUTHORISATION

We apply zero trust concepts [NIST SP 800-207] to agent systems with four stipulations:

• Per-call authorisation: Each tool invocation is independently sanctioned based on the convergence of agent capabilities, user permissions, and task parameters. Evaluation at the action level rather than the session level.

• Identity-inherited permissions: Agents get the scoped permissions of the initiating user rather than ambient credentials. This confines the blast radius of any breach to the actions that the initiating user may have personally executed.

• Sequence-aware policy evaluation: Analysing action sequences to identify composition assaults. New activities are assessed in relation to previous acts inside the session, rather than in isolation. This necessitates the preservation of activity history and the implementation of pattern recognition for illicit compositions.

• Continuous re-evaluation: Authorisation is reassessed whenever the context alters during execution (user permissions are revoked, data sensitivity is reclassified, or time-based policy modifications occur).

ARCHITECTURE OF DELEGATION TOKENS

We suggest the enhancement of OAuth 2.0 token exchange (RFC 8693) for the purpose of agent delegation:

User activates agent using authorised identity (SSO, OIDC token). Agent runtime

generates a delegation token that encodes: Agent Y representing User X, for task Z, within scope W, expiring at a specified period.

When an agent invokes a tool or resource, the delegation token is sent in conjunction with the agent's own identity.

The resource assesses both: Is the agent permitted to make this call? Is User X permitted within this scope? Both must be approved for execution.

In multi-agent systems, delegation chains are upheld by agent-to-agent transfers, with each connection cryptographically authenticated and the initial human identity transmitted across the entire chain.

ENFORCEMENT AT THE INFRASTRUCTURE LEVEL

We differentiate between two essentially distinct governance models:

Trust-based governance assigns policy directives within prompts or configuration files that the model interprets and elects to adhere to. This model offers no security assurances and is susceptible to rapid injection, model inaccuracies, jailbreaking, and the intrinsic helpfulness-safety tradeoff seen in RLHF-trained models.

Enforcement-based governance situates policy evaluation at an infrastructural level that the agent is unable to access or alter. Authorisation decisions are rendered by deterministic systems beyond the model's contextual scope. The policy enforcement layer intercepts tool calls prior to execution, assessing them against authorisation policies utilising the delegation token, capability manifest, and action history. This model alone offers the requisite assurance level for enterprise deployment.

VERIFICATION OF INDEPENDENT ACTION

To resolve the oracle dilemma, we advocate for independent verification of outcomes reported by agents. Essential activities are validated by directly asking target systems to ensure the activity was performed as indicated. The verification layer functions autonomously from the agent, contrasting agent-reported results with the real system state. This is comparable to independent verification in pharmaceutical manufacture, where each production phase necessitates validation by a second trained observer.

AUDIT AWARE OF HALLUCINATIONS

Conventional audit methods presume that recorded actions accurately represent real occurrences. In agent systems, this assumption is invalid due to: (a) the agent potentially fabricating acts it did not perform, (b) the agent perhaps conjuring parameters for actions it did execute, and (c) chain-of-thought reasoning sequences may serve as retrospective justifications. We propose audit systems that: document both agent-reported actions and independently validated outcomes; identify discrepancies between reported and validated actions; regard reasoning chains as ancillary evidence rather than

definitive records; and correlate actions across agents in multi-agent systems to uncover patterns not discernible at the individual agent level.

EMPIRICAL EVIDENCE FROM PRODUCTION SYSTEMS

CLAWHAVOC SUPPLY CHAIN ASSAULT

The ClawHavoc campaign (CVE-2026-25253) signifies the inaugural recorded extensive supply chain assault aimed at an AI agent ecosystem. Koi Security detected over 824 harmful skills in the OpenClaw ClawHub marketplace, comprising infostealers intended to extract saved credentials from recognised file paths, keyloggers that record user input during agent sessions, and backdoors that provide ongoing remote access. A singular concerted campaign was responsible for 335 of these nefarious packages. OpenClaw's design retains API keys, conversation history, and webhook tokens in predictable local file directories, establishing concentrated targets for such assaults.

ANALYSIS OF SKILL VULNERABILITY

Cisco's security assessment of 31,000 OpenClaw agent skills revealed that 26% exhibited vulnerabilities. A widely utilised skill harboured 2 critical and 5 high-severity vulnerabilities that facilitated covert data exfiltration through embedded curl instructions. OpenClaw's security literature concedes that "prompt injection remains unresolved" and that "system prompt guardrails are merely soft guidance" - a notable transparency that highlights the disparity between existing defences and the necessary confidence levels.

MULTI-AGENT CREDENTIAL EVALUATION

Our examination of credential transmission throughout three prominent multi-agent frameworks uncovers a uniform trend of ambient authority inheritance:

• Claude Code Agent Teams: All subordinate agents adopt the authorisation settings of the main agent. There is no per-agent scope inside the team.

• Cursor FastRender: 2,000 simultaneous agents with hierarchical role differentiation yet consistent credential scope across all roles.

• OpenClaw: More than 30,000 instances are accessible on the public internet. Uniform credential framework devoid of identity delegation.

In none of these systems can agent actions be attributed to a specific initiating user, nor can permissions be delineated according to the user's authorisation level.

DISCUSSION

INFRASTRUCTURE ENFORCEMENT: ESSENTIAL YET INSUFFICIENT

The central tenet of our system — that security assurances necessitate infrastructure-

level regulation beyond the model's jurisdiction — mitigates most authorisation threats. We recognise the constraints. Composition attacks illustrate that per-call infrastructure enforcement can be evaded through sequences of separately authorised activities. Sequence-aware policy assessment partially mitigates this issue but cannot predict all novel compositions that emergent model reasoning may uncover. Layered defence the integration of infrastructure enforcement, sequence analysis, behavioural anomaly detection, and human oversight for high-risk operations constitutes the most comprehensive strategy.

THE MODEL AS OPPONENT

A security framework that regards the LLM as a compliant and predictable element will consistently underestimate danger. The essential characteristics of huge language models — non-determinism, hallucination, adversary susceptibility, alignment instability, and emergent capabilities — should be regarded as primary security concerns. This does not imply the abandonment of model-level controls; rather, it recognises that they offer probabilistic risk mitigation, but infrastructure-level controls ensure deterministic enforcement.

CRYPTOGRAPHIC GOVERNANCE LACKING DECENTRALISATION

Our platform utilises cryptographic primitives—signed delegation tokens, cryptographic agent identity, immutable audit logs, and verifiable action chains—without necessitating decentralised consensus. Enterprise demands for centralised policy, immediate revocation, minimal latency, and data privacy are incongruent with blockchain-based governance. We utilise cryptography tools from decentralised systems while preserving centralised policy authority. Cross-organizational agent interactions necessitate federated identity models where neither organization governs the other's identity infrastructure.

CONSTRAINTS

Our framework has not undergone validation in large-scale production deployments. The performance overhead of per-call authorisation with sequence analysis in high-throughput multi-agent systems, which involve thousands of tool calls per hour, necessitates empirical evaluation. The independent action verification mechanism introduces latency and complexity that may be unfeasible for all actions, necessitating risk-based tiering. Ultimately, our composition assault detection depends on pattern matching with established attack patterns; innovative compositions may circumvent detection until patterns are revised.

REFERENCES

[1] Thalary, S., & Katipelly, A. (2021). CI/CD for Distributed Software Systems: Why Software Architecture Determines Pipeline Complexity. *International Journal of Emerging Research in Engineering and*

Technology, *2*(4), 100-111.

[2] Fletcher, L., & Lysova, E. I. (2026). Addressing gender in authenticity and inclusion at work: Nuancing conservation of resources theory with social role theory. *Human Resource Management*, *65*(1), 219-234.

[3] Kuntamukkala, N. K., & Katipelly, A. (2022). Neural Component Libraries for Angular: AI-Generated, Self-Documenting UI Elements with Intelligent API Integration. *International Journal of AI, BigData, Computational and Management Studies*, *3*(3), 116-127.

[4] Mukherjee, S. (2026). in Adolescents: An Eclectic Clinical Framework. *Applied Mindfulness and Integrative Practices for Mental Well-Being*, 159.

[5] Katipelly, A. (2022). Hierarchical Multi-Agent Orchestration for Automated Dispute Resolution. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, *3*(3), 140-150.

[6] Wadehra, S., Chatterjee, R., & Dubey, S. (2026). Instruction Manual: Code of silence: Speak Stronger's mission beyond borders. *Teaching Notes*, 1-56.

[7] Katipelly, A., & Kuntamukkala, N. K. (2022). Mitigating Algorithmic Complexity Attacks in Federated GraphQL Architectures: A Depth-Bounded Semantic Rate Limiting Approach for Open Banking. *International Journal of Emerging Trends in Computer Science and Information Technology*, *3*(3), 112-121.

[8] Fletcher, L., & Lysova, E. I. (2026). Addressing gender in authenticity and inclusion at work: Nuancing conservation of resources theory with social role theory. *Human Resource Management*, *65*(1), 219-234.

[9] Katipelly, A., & Thalary, S. (2023). Cryptographic Identity Propagation in Asynchronous Event-Driven Architectures: Implementing Zero-Trust Envelopes for High-Velocity Payment Streams. *International Journal of Emerging Trends in Computer Science and Information Technology*, *4*(2), 212-222.

[10] Bhambri, P., & Hamad, A. (2025). Quantum computing and AI synergies: Strengthening cybersecurity resilience. In *Handbook of AI-Driven Threat Detection and Prevention* (pp. 337-352). CRC Press.

[11] Kuntamukkala, N. K., & Katipelly, A. (2023). Predictive Angular Rendering: Machine Learning Models for Intelligent Client-Side Optimization with Adaptive Backend Coordination. *International Journal of AI, BigData, Computational and Management Studies*, *4*(2), 144-154.

[12] Doria, B., Grion, V., Picasso, F., & Li, L. (2026). Faculty assessment development in higher education: the CRAFT framework emerging from a systematic literature review. *Assessment & Evaluation in Higher Education*, *51*(1), 16-40.

[13] Thalary, S., & Katipelly, A. (2023). Secure-by-Design Cloud Software Delivery: How DevOps and Software Teams Co-Own Security Outcomes. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, *4*(1), 131-140.

[14] Katipelly, A. (2024). Hierarchical Agentic Orchestration for Microservices: A Neuro-Symbolic Framework for Dynamic Workflow Composition in Decentralized Financial Systems. *International Journal of Emerging Research in Engineering and Technology*, *5*(4), 165-174.

[15] Perez, M. V., Linley, J., Stroup, N. R., Tennessen, N. F., & Gyesi, M. (2026). "My Biggest Thing is Helping People, Helping Others": A Qualitative Study of Peer Socialization Agents' Motivations. *Journal of Student Affairs Research and Practice*, *63*(1), 117-129.

[16] Katipelly, A., & Thalary, S. (2024). Semantic Automation of Basel III Liquidity Reporting: Utilizing Ontological Knowledge Graphs for Real-Time Regulatory Compliance and Auditability. *International Journal of Emerging Research in Engineering and Technology*, *5*(2), 147-156.

[17] Wadehra, S., Chatterjee, R., & Dubey, S. (2026). Instruction Manual: Code of silence: Speak Stronger's mission beyond borders. *Teaching Notes*, 1-56.

[18] Thalary, S., & Katipelly, A. (2024). Cloud-Native Design for Event-Driven Systems: Where Software Architecture Decisions Meet DevOps Reality. *International Journal of AI, BigData, Computational and Management Studies*, *5*(2), 202-212.

[19] Perez, M. V., Linley, J., Stroup, N. R., Tennessen, N. F., & Gyesi, M. (2026). "My Biggest Thing is Helping People, Helping Others": A Qualitative Study of Peer Socialization Agents' Motivations. *Journal of Student Affairs Research and Practice*, *63*(1), 117-129.

[20] Katipelly, A. (2024). Predictive AI Proactive Customer Engagement Platform and Real-Time Friction Reduction Using AI-Based Churn Prediction. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, *5*(1), 211-221.

[21] Wadehra, S., Chatterjee, R., & Dubey, S. (2026). Instruction Manual: Code of silence: Speak Stronger's mission

beyond borders. *Teaching Notes*, 1-56.

[22] Katipelly, A., & Thalary, S. (2025). Carbon-Aware Dynamic Batching for Deep Learning Inference: Optimizing the Energy-Latency Trade-off in High-Frequency Transaction Monitoring. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, *6*(3), 160-169.

[23] Thinn, M. (2026). *Kantian Ethics and Principled Leadership in Profit-Driven Organizations* (Doctoral dissertation, South College).

[24] Katipelly, A., & Kuntamukkala, N. K. (2025). Hierarchical Multi-Agent Orchestration for Automated Dispute Resolution: A Game-Theoretic Approach to Policy Adherence in Digital Wallets. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, *6*(2), 195-204.

[25] Katsaros, K., Malisova, O., Lazanaki, V., & Tsoni, E. (2026). Stirring inclusion: how diversity-oriented HR practices boost adaptive performance in Greece's food & beverage industry. *Frontiers in Psychology*, *17*, 1768113.

[26] Thalary, S., & Katipelly, A. (2025). Platform Engineering for Distributed Systems: How Cloud DevOps Enables Scalable, Policy-Driven Software Architectures. *International Journal of AI, BigData, Computational and Management Studies*, *6*(1), 189-197.

[27] Murtagh, L., Parker, K., & Sullivan, S. (2026). Mapping critical moments: a collaborative autoethnographical exploration of social justice formation in leaders of teacher educator programmes using life grid methodology. *Journal of Education for Teaching*, 1-15.

[28] Katipelly, A., Sethuraman, P., & Manoj, M. S. (2026, March). An AI-Powered Event-Driven Architecture for Predictive Customer Experience Management Using Autonomous Decision Agents. In *2026 IEEE Madhya Pradesh Section Conference (MPCON)* (pp. 1084-1091). IEEE.

[29] Sethuraman, P., Katipelly, A., & Manoj, M. S. (2026, March). Event-Driven Reinforcement Learning Frameworks for Proactive Customer Friction Detection in Large-Scale Digital Platforms. In *2026 IEEE Madhya Pradesh Section Conference (MPCON)* (pp. 1686-1693). IEEE.