

# Zero-Trust Identity Fabric for Autonomous AI Agents: Design, Security Controls, and Empirical Evaluation

Haruto Nakamura

Artificial Intelligence Researcher Japan

Yui Tanaka

Artificial Intelligence Research Analyst Japan

VOLUME 03 ISSUE 09 (2026)

Published Date: 10 September 2026 // Page no.: - 08-14

---

## ABSTRACT

The increasing autonomy of artificial intelligence agents introduces a fundamental identity-security problem: conventional identity architectures generally authenticate human users and relatively static applications, whereas autonomous agents can dynamically invoke tools, exchange information, delegate subtasks, and make decisions with limited direct human intervention. This paper proposes a Zero-Trust Identity Fabric (ZTIF) for autonomous AI agents that treats every agent, delegated action, tool invocation, and inter-agent interaction as an independently verifiable security event. The proposed architecture integrates continuous identity verification, contextual authorization, least-privilege access, cryptographically protected delegation, policy enforcement, behavioral monitoring, and auditable decision trails. The methodological foundation combines identity-fabric principles with concepts derived from adaptive control, optimization, simulation, data mining, fault diagnosis, and safe planning represented in the supplied literature. In particular, the identity-fabric perspective is directly informed by Pappu, Bhushan, and Jaiswal's work on future-proofing identity security for agentic AI systems (Pappu et al., 2026). The paper develops a layered architectural model and evaluates it through a controlled conceptual test framework covering identity compromise, privilege escalation, unauthorized tool invocation, delegation abuse, and anomalous agent behavior. The evaluation indicates that zero-trust controls are most effective when identity, authorization, behavioral context, and action-level policy are evaluated jointly rather than independently. The study contributes an integrated security model for autonomous AI environments and identifies important trade-offs between security assurance, operational latency, policy complexity, and agent autonomy.

**Keywords:** Zero Trust; Identity Fabric; Autonomous AI Agents; Agentic AI; Identity Security; Continuous Authentication; Adaptive Authorization; Privileged Access; AI Security; Policy Enforcement.

---

## INTRODUCTION

Autonomous AI agents represent a transition from software systems that primarily respond to explicit user commands toward systems capable of planning, selecting tools, executing workflows, and adapting their behavior according to intermediate results. This transformation creates a distinctive identity problem. An agent is not merely a software process: during execution, it may assume delegated authority, access multiple resources, interact with other agents, and initiate actions whose consequences extend beyond the original user request. Consequently, authentication performed only at session initiation is insufficient for environments in which authorization conditions can change throughout an agent's lifecycle.

Traditional identity architectures commonly emphasize a relatively stable relationship among a human identity, an application, and a resource. Autonomous agents disrupt this relationship because an agent may operate

simultaneously as an executor, delegator, consumer, and producer of information. A security architecture must therefore establish not only who an agent is but also what authority it currently possesses, why that authority was granted, what action it is attempting, and whether the action remains consistent with applicable policy.

This problem is particularly important because increased autonomy creates a larger operational attack surface. An attacker who compromises an agent's credentials may potentially exploit the agent's legitimate permissions rather than directly bypassing authentication. Similarly, an agent that has been legitimately authorized to access one system could misuse that authority through an unexpected sequence of tool calls. Such situations demonstrate why identity assurance and authorization cannot be treated as isolated controls.

The present study addresses this problem through a Zero-Trust Identity Fabric (ZTIF), in which identity verification and authorization are continuously applied to agents, users, services, tools, delegated credentials, and individual high-risk

actions. The architectural perspective is aligned with the identity-fabric approach proposed for agentic AI security by Pappu, Bhushan, and Jaiswal, which specifically positions identity security as a foundational requirement for future agentic systems (Pappu et al., 2026).

The primary objective is to design a security architecture capable of maintaining strong identity assurance without eliminating useful agent autonomy. The study therefore investigates three questions. First, what architectural components are required to establish trustworthy identities for autonomous agents? Second, how can authorization be dynamically evaluated when agent context and behavior change during execution? Third, how can the security effectiveness of such an architecture be evaluated against representative agent-specific threats?

The significance of this research extends beyond authentication. Autonomous systems require security controls capable of connecting identity, context, behavior, authorization, and accountability. The proposed architecture consequently treats an agent action as a security decision rather than merely an execution event.

## LITERATURE REVIEW

The supplied literature does not constitute a conventional identity-security corpus; instead, it contains studies spanning data mining, adaptive control, robotics, optimization, neural networks, simulation, fault diagnosis, and mathematical modeling. Their relevance to this study is therefore methodological rather than domain-specific. These works collectively provide useful foundations for understanding adaptive decision-making, anomaly identification, optimization, simulation, and safety constraints that can be transferred conceptually to autonomous-agent identity management.

Chen, Han, and Yu provide an important foundation through their overview of data mining from a database perspective. Their work emphasizes systematic extraction of meaningful patterns from data (Chen et al., 1996). For an identity fabric, this principle is relevant because continuous authorization can incorporate behavioral evidence rather than relying exclusively on static identity attributes. An agent repeatedly accessing an unusual resource, invoking an unexpected tool, or deviating from established interaction patterns can generate contextual evidence for an authorization engine.

Fault-diagnosis research provides another useful conceptual connection. Chao et al. combine deep neural networks with system-performance models for fault detection and isolation (Chao et al., 2019). The underlying principle is significant for agent security: security monitoring can benefit from combining learned

behavioral signals with explicit system expectations. Instead of simply detecting whether an agent's behavior is unusual, a security mechanism can evaluate whether the behavior is inconsistent with the agent's authorized operational model.

Sadoughi and Hu similarly investigate physics-based convolutional neural networks for fault diagnosis (Sadoughi & Hu, 2019). Although their application concerns rolling-element bearings rather than cybersecurity, their work demonstrates the broader value of incorporating structured system knowledge into learning-based detection. In an identity fabric, behavioral monitoring should likewise combine statistical signals with explicit policy constraints.

Adaptive control research provides another theoretical analogy. Kamalasadan and Ghandakly examine a neural-network-based parallel adaptive controller for dynamic systems (Kamalasadan & Ghandakly, 2007). Adaptive control is relevant because autonomous agents operate in changing environments. A fixed authorization decision may become inappropriate when the agent's context changes. However, adaptive authorization must be constrained: unlike unrestricted learning, security adaptation should not permit the system to autonomously expand its own privileges.

Optimization-oriented studies further support this perspective. Fang, Xin, and Zhang use interval analysis, region mapping, and differential evolution to address constrained robotic design (Fang et al., 2015), while Jamwal et al. examine performance-based optimization of a compliant parallel robot (Jamwal et al., 2022). These studies demonstrate the importance of optimizing system performance subject to explicit constraints. The same principle applies to identity fabrics: authorization decisions should optimize operational utility while remaining inside security boundaries.

Integrated design is particularly relevant to the proposed architecture. Alandoli et al. investigate dynamic modeling and integrated optimal control of hybrid robotic arms (Alandoli et al., 2023), while Martínez-Castelán and Villarreal-Cervantes examine integrated structure-control design for a bipedal robot (Martínez-Castelán & Villarreal-Cervantes, 2021). These studies support an architectural argument that security should not be attached as an isolated layer after agent functionality has been designed. Identity, policy, execution, and monitoring should be integrated into the agent lifecycle.

Kober, Bagnell, and Peters provide a survey of reinforcement learning in robotics (Kober et al., 2013). Reinforcement-learning concepts are relevant to autonomous agents because agents can select actions dynamically. From a security perspective, however, learned action selection introduces a critical distinction between decision autonomy and security authority. An agent may be permitted to choose among actions without being permitted to expand its authorization scope.

Simulation provides another useful methodological

foundation. Lee et al. investigate privileged adversarial learning from simulation (Lee et al., 2018). The relevance here is the use of simulated environments to expose systems to conditions that may be difficult or unsafe to reproduce operationally. A zero-trust identity architecture can similarly be evaluated through simulated agent compromise, abnormal delegation, and unauthorized tool-use scenarios.

Safe planning research provides a particularly strong conceptual analogy. Yakovlev examines safe interval path planning and flatness-based control for navigation among static and dynamic obstacles (Yakovlev, 2022). Security authorization can be viewed similarly as constrained action planning: an agent may have multiple possible actions, but only a subset is permissible under the current security state.

Palacios's mathematical modeling approach provides a broader theoretical basis for representing complex systems through formal models (Palacios, 2022). For identity fabrics, formal modeling can describe relationships among identities, privileges, contexts, policies, resources, and actions.

Raissi, Perdikaris, and Karniadakis demonstrate data-driven solutions for nonlinear partial differential equations through physics-informed deep learning (Raissi et al., 2017). Although the application differs substantially, the work reinforces the possibility of combining learned information with explicit structural constraints. Such a hybrid principle is valuable for security systems because purely learned security decisions may be difficult to explain or constrain.

Blechschiidt and Ernst's review of neural-network approaches to partial differential equations emphasizes that multiple methodological approaches can exist for solving complex computational problems (Blechschiidt & Ernst, 2021). For agent identity security, this suggests that no single detection or authorization mechanism should be assumed universally sufficient.

The supplied literature therefore reveals a significant gap. Existing references provide concepts for adaptive control, optimization, anomaly detection, simulation, and constrained decision-making, while the directly relevant identity-fabric contribution is represented by Pappu et al. (2026). What remains insufficiently integrated is an architecture that combines these ideas into an action-level zero-trust identity model specifically designed for autonomous AI agents. This paper addresses that gap by treating identity, context, policy, behavior, delegation, and action as interconnected security variables.

## METHODOLOGY

### 3.1 Research Design

This study adopts a design-science-oriented methodology consisting of architectural modeling, security-control specification, threat-oriented scenario construction, and empirical-style evaluation criteria. The objective is not to treat an agent as a conventional authenticated application but to model its identity as a continuously evaluated security state.

The proposed ZTIF is represented conceptually as:

Agent Identity → Context Verification → Policy Evaluation → Authorization Decision → Action Execution → Behavioral Observation → Trust Reassessment

This feedback cycle distinguishes the proposed architecture from static access-control models. Every significant action generates evidence that can influence subsequent decisions.

The methodological approach is consistent with the integrated-design logic observed in the supplied literature, where system performance and constraints are considered jointly rather than independently (Alandoli et al., 2023; Jamwal et al., 2022).

### 3.2 Identity Model

Each autonomous agent receives a cryptographically verifiable identity containing a unique identifier, credential metadata, ownership or provenance information, permitted capabilities, lifecycle status, and delegation constraints.

The identity fabric distinguishes at least four identity relationships: human-to-agent, agent-to-agent, agent-to-tool, and agent-to-resource. This distinction is necessary because successful authentication of an agent does not automatically establish authorization for every downstream operation.

For example, an analytical agent may be authorized to read financial data but prohibited from modifying records. If it delegates a task to another agent, the delegated agent should not inherit broader privileges than those available to the originating agent.

This establishes a fundamental principle:

Delegation must preserve or reduce privilege; it must never silently amplify privilege.

The identity-fabric concept is particularly important here because future agentic environments require identity controls that remain meaningful across dynamically changing agent relationships (Pappu et al., 2026).

### 3.3 Context-Aware Authorization

The proposed authorization engine evaluates more than identity. A decision is generated from a contextual tuple:

$$A = f(I, R, X, P, B, D, T)$$

where:

- I = verified agent identity,
- R = requested resource,
- X = current execution context,
- P = applicable security policy,
- B = observed behavior,
- D = delegation state,
- T = temporal and trust conditions.

The resulting decision can be expressed as Permit, Deny, or Step-Up Verification.

Step-up verification is particularly useful for autonomous agents. Rather than terminating every unusual operation, the identity fabric can require additional evidence before permitting a sensitive action.

For example, an agent may normally retrieve customer information but suddenly request bulk deletion. The identity remains valid, but the action falls outside the agent's normal operational profile. A contextual authorization engine can therefore require elevated approval or deny the operation.

### 3.4 Least-Privilege Capability Model

The architecture uses capability-oriented authorization to reduce excessive privileges. An agent should receive narrowly defined capabilities rather than broad access to an entire application environment.

Capabilities can be constrained by resource, action type, time, transaction value, and delegation depth. A capability may therefore allow an agent to read a specific dataset for a limited execution period while prohibiting export, modification, or redistribution.

Optimization literature provides an appropriate conceptual analogy because constrained optimization seeks useful outcomes while maintaining defined boundaries (Fang et al., 2015). In the identity fabric, operational efficiency cannot supersede security constraints.

### 3.5 Continuous Behavioral Monitoring

Static credentials cannot reliably distinguish legitimate use from compromised use. The proposed architecture therefore maintains a behavioral profile consisting of action frequency, resource selection, tool sequences, delegation patterns, execution timing, and error behavior.

Data-mining principles are relevant because large volumes of event information can be transformed into behavioral patterns (Chen et al., 1996). Likewise, fault-detection research demonstrates the usefulness of combining observed behavior with system models (Chao et al., 2019).

A practical detection mechanism could calculate an anomaly score:

$$S = w_1I + w_2C + w_3B + w_4D + w_5R$$

where the terms represent identity confidence, contextual deviation, behavioral anomaly, delegation risk, and resource sensitivity.

A high score does not necessarily imply malicious behavior. Instead, it increases the security requirements associated with the next action.

### 3.6 Delegation Security

Delegation is one of the most significant security differences between autonomous agents and conventional applications. An agent may ask another agent to complete a task, producing a chain of authority.

The proposed architecture records:

Originator → Delegator → Delegatee → Tool → Resource → Action

Each transition must preserve provenance. A downstream agent must be able to demonstrate where its authority originated and what limitations accompanied that authority.

Delegation depth should also be bounded. For example, an agent authorized to generate a report may delegate data retrieval to another agent but should not permit that agent to create additional unrestricted delegates.

### 3.7 Policy Enforcement Point

All sensitive actions pass through a policy enforcement point. This component separates agent intention from authorized execution. The agent may request an action, but the enforcement point determines whether that action is permitted.

This separation is essential because autonomous systems can generate actions dynamically. The security boundary must therefore exist outside the agent's own reasoning process.

The principle resembles safe planning under constraints: multiple actions may be technically possible, but only some belong to the permissible operational region (Yakovlev, 2022).

### 3.8 Evaluation Framework

The empirical evaluation framework uses controlled scenarios

representing common identity-security failures:

1. compromised agent credential;
2. unauthorized privilege escalation;
3. abnormal tool invocation;
4. excessive delegation;
5. anomalous resource access;
6. legitimate but high-risk action;
7. policy violation caused by contextual change.

Each scenario is evaluated against baseline static authentication and the proposed ZTIF. The primary evaluation dimensions are unauthorized-action prevention, detection responsiveness, policy precision, delegation containment, and operational overhead.

Simulation-based evaluation is appropriate because high-risk agent compromises may not be safely reproduced in production environments. The role of simulated environments in evaluating intelligent systems is conceptually supported by the simulation-oriented approach represented by Lee et al. (2018).

## RESULTS

The evaluation framework indicates that the proposed ZTIF provides stronger security coverage when identity verification, contextual authorization, and behavioral monitoring operate as a unified control loop. Under a conventional static-authentication model, an agent possessing valid credentials can potentially continue executing actions after its operational context has changed. The ZTIF instead reassesses authorization at action boundaries, reducing the dependence on the initial authentication event.

The first major finding concerns credential compromise. Possession of a valid agent identity is insufficient for unrestricted activity because the proposed model evaluates the requested resource, action type, behavioral profile, and current policy state. Consequently, a compromised credential can remain syntactically valid while still producing authorization failures when its behavior deviates substantially from the agent's established operating profile.

The second finding concerns privilege escalation. Capability-based authorization and delegation constraints prevent a downstream agent from automatically inheriting the complete authority of an upstream identity. The architecture therefore transforms privilege inheritance from an implicit property into an explicit

security decision. This is particularly important in multi-agent workflows, where privilege amplification can occur indirectly through legitimate delegation.

The third finding concerns behavioral anomalies. A monitoring layer based on data-derived patterns can identify deviations that static identity verification cannot detect. This finding is consistent with the methodological value of data mining for identifying meaningful patterns in large datasets (Chen et al., 1996). However, behavioral evidence is most effective when combined with explicit policy constraints rather than used as the sole authorization mechanism.

The fourth finding concerns high-risk legitimate actions. The architecture distinguishes between an invalid identity and a valid identity performing an unusually sensitive operation. This permits step-up verification rather than indiscriminate denial. Such adaptive control improves operational flexibility while preserving a security boundary, reflecting the broader value of adaptive decision mechanisms demonstrated in control-oriented research (Kamalasadan & Ghandakly, 2007).

The fifth finding concerns delegation containment. Explicit provenance and bounded delegation reduce the possibility that an agent can construct an uncontrolled chain of authority. The security model consequently changes the fundamental question from "Is this agent authenticated?" to "Is this exact action authorized through a valid chain of authority?"

Finally, the evaluation identifies a trade-off between security assurance and execution latency. Continuous authorization, behavioral analysis, and policy evaluation introduce computational and operational overhead. The results therefore indicate that zero-trust security should use risk-sensitive enforcement: low-risk operations can use lightweight checks, whereas sensitive actions require stronger contextual evaluation.

Overall, the findings support the central proposition that an autonomous AI identity architecture should be action-centric rather than session-centric. Identity remains necessary, but identity alone is insufficient for trustworthy autonomy.

## DISCUSSION

The findings have important theoretical implications for the development of autonomous AI security. Conventional identity models frequently treat identity as a relatively stable attribute, whereas autonomous agents require identity to function as a continuously evaluated security state. The proposed ZTIF therefore extends the identity concept from authentication toward an integrated relationship among identity, authority, context, behavior, and action. This direction is consistent with the identity-fabric perspective developed specifically for agentic AI systems by Pappu et al. (2026).

The first important implication is that trust should not be transferred automatically across agent interactions. In human-

centered systems, delegation may sometimes be modeled through organizational roles. Autonomous agents, however, can create delegation chains at machine speed. A trusted originating agent can therefore unintentionally become a mechanism for distributing excessive authority. The proposed bounded-delegation mechanism addresses this issue by requiring each delegation relationship to preserve provenance and privilege constraints.

The second implication concerns the role of machine learning. The literature demonstrates substantial value in neural networks, adaptive control, and data-driven modeling (Blechsmidt & Ernst, 2021; Raissi et al., 2017). Nevertheless, applying purely learned decision-making to authorization would introduce unacceptable uncertainty in high-risk environments. Security authorization should therefore follow a hybrid architecture in which machine learning contributes behavioral evidence while deterministic policies establish non-negotiable security boundaries. This resembles the broader methodological principle of combining learned representations with explicit system knowledge found in fault-diagnosis research (Sadoughi & Hu, 2019).

A third implication involves optimization. Autonomous agents are designed to maximize task completion and operational efficiency, while security systems are designed to constrain harmful behavior. These objectives can conflict. An agent may identify a faster route to completing a task by requesting broader permissions, but a zero-trust architecture should reject such implicit privilege expansion. Constrained optimization concepts in engineering research reinforce the principle that performance must be evaluated inside defined operational boundaries (Fang et al., 2015; Jamwal et al., 2022).

The principal practical limitation is policy complexity. As the number of agents, tools, resources, and delegation relationships increases, policy management can become difficult. Excessively restrictive policies can also reduce the usefulness of autonomous systems. A second limitation is false-positive behavioral detection. Legitimate agents may change their behavior when assigned novel tasks, potentially triggering anomaly mechanisms. A third limitation is computational overhead. Continuous authorization and monitoring require additional processing, especially when agents execute large numbers of actions.

Another limitation concerns the empirical status of the present evaluation. The framework specifies controlled evaluation scenarios and analytical findings but does not claim measurements from a production deployment or a supplied benchmark dataset. Therefore, the findings should be interpreted as architecture-level evaluation outcomes rather than definitive statistical estimates of

security improvement. Future experimental work should implement the model in a multi-agent testbed and measure detection accuracy, latency, false-positive rates, policy decision throughput, and attack containment.

Despite these limitations, the architecture provides a coherent foundation for agentic identity security because it separates autonomy from authority. Agents can remain autonomous in selecting actions while security infrastructure independently determines whether those actions are permissible. This separation is essential for building systems that are both operationally useful and resistant to identity-based compromise.

## CONCLUSION

This paper proposed a Zero-Trust Identity Fabric for autonomous AI agents that integrates continuous identity verification, contextual authorization, least-privilege capabilities, behavioral monitoring, secure delegation, policy enforcement, and action-level auditing. The central contribution is a shift from session-centric identity toward action-centric trust evaluation.

The analysis demonstrates that autonomous-agent security cannot depend exclusively on credential validity. An authenticated agent may still perform an unauthorized action, misuse delegated privileges, invoke an inappropriate tool, or behave anomalously after its initial authentication. Consequently, identity must be continuously evaluated in relation to context, policy, behavior, delegation, and resource sensitivity.

The provided literature supports the methodological foundations of this approach through research on data mining, adaptive control, constrained optimization, fault diagnosis, simulation, safe planning, and integrated system design. Most importantly, the identity-fabric perspective directly supports the need for future-proof identity security in agentic AI environments (Pappu et al., 2026).

The proposed architecture also establishes a practical separation between agent autonomy and security authority. Agents can make dynamic decisions, but they cannot unilaterally determine their own permissions. This principle provides a scalable foundation for multi-agent environments in which identity relationships and delegated actions may become highly complex.

Future research should implement the proposed model in real multi-agent platforms and conduct quantitative experiments involving adversarial agents, compromised credentials, dynamic delegation, tool misuse, and policy conflicts. Additional research should investigate cryptographic delegation, explainable authorization decisions, policy-learning mechanisms, and low-latency enforcement architectures. Such developments can help transform zero-trust identity from a conventional enterprise security concept

into a foundational security architecture for autonomous artificial intelligence.

## REFERENCES

1. M. A. Chao, C. Kulkarni, K. Goebel, and O. Fink, "Hybrid deep fault detection and isolation: Combining deep neural networks and system performance models," 2019, arXiv:1908.01529.
2. M. S. Chen, J. Han, and P. S. Yu, "Data mining: An overview from a database perspective," *IEEE Trans. Knowl. Data Eng.*, vol. 8, no. 6, pp. 866–883, Dec. 1996.
3. H. Fang, B. Xin, and X. Zhang, "Workspace-constrained optimal design of three-degrees-of-freedom parallel manipulators with minimum parasitic motions by integrating interval analysis, region mapping, and differential evolution," *Eng. Optim.*, vol. 47, no. 3, pp. 407–428, 2015, doi: 10.1080/0305215X.2014.895337.
4. P. K. Jamwal, A. Kapsalyamov, S. Hussain, and M. H. Ghayesh, "Performance-based design optimization of an intrinsically compliant 6-DOF parallel robot," *Mech. Based Des. Struct. Mach.*, vol. 50, no. 4, pp. 1237–1252, 2022, doi: 10.1080/15397734.2020.1746669.
5. S. Kamalasadan and A. A. Ghandakly, "A neural network parallel adaptive controller for dynamic system control," *IEEE Trans. Instrum. Meas.*, vol. 56, no. 5, pp. 1786–1796, Oct. 2007, doi: 10.1109/TIM.2007.895674.
6. J. Kober, J. A. Bagnell, and J. Peters, "Reinforcement learning in robotics: A survey," *Int. J. Robot. Res.*, vol. 32, no. 11, pp. 1238–1274, 2013.
7. K.-H. Lee, G. Ros, J. Li, and A. Gaidon, "SPIGAN: Privileged adversarial learning from simulation," 2018, arXiv:1810.03756.
8. J. N. Martínez-Castelán and M. G. Villarreal-Cervantes, "Integrated structure-control design of a bipedal robot based on passive dynamic walking," *Mathematics*, vol. 9, no. 13, p. 1482, 2021, doi: 10.3390/math9131482.
9. A. Palacios, *Mathematical Modeling: A Dynamical Systems Approach to Analyze Practical Problems in STEM Disciplines*. Cham, Switzerland : Springer International Publishing, 2022.
10. M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics informed deep learning (part I): Data-driven solutions of nonlinear partial differential equations," 2017, arXiv:1711.10561.
11. M. Sadoughi and C. Hu, "Physics-based convolutional neural network for fault diagnosis of rolling element bearings," *IEEE Sens. J.*, vol. 19, no. 11, pp. 4181–4192, Jun. 2019.
12. E. A. Alandoli, T. S. Lee, V. Vijayakumar, Y. J. Lin, and M. Q. Mohammed, "Dynamic model and integrated optimal controller of hybrid arms robot for laser contour machining," *J. Vib. Control*, vol. 29, nos. 13–14, pp. 3006–3024, 2023, doi: 10.1177/10775463221090000.
13. K. S. Yakovlev, "Safe interval path planning and flatness-based control for navigation of a mobile robot among static and dynamic obstacles," *Autom. Remote Control*, vol. 83, pp. 903–918, 2022, doi: 10.1134/S000511792206008X.
14. J. Blechschmidt and O. G. Ernst, "Three ways to solve partial differential equations with neural networks—A review," *GAMM-Mitteilungen*, vol. 44, no. 2, 2021, Art. no. e202100006, doi: 10.1002/gamm.202100006.
15. K. Pappu, B. Bhushan and N. Jaiswal, "Future-Proofing Identity Security for Agentic AI Systems: Design, Implementation, and Evaluation of an Identity Fabric," 2026 International Conference on Artificial Intelligence, Systems, and Emerging Technologies (ICAISSET), Cairo, Egypt, 2026, pp. 1-7, doi: 10.1109/ICAISSET66439.2026.11541783.