International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 05 ISSUE 10 Pages: 87-96
OCLC – 1368736135

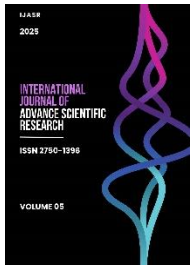

Journal Website: http://sciencebring.com/index.php/ijasr

Research Article

An Explainable Zero-Trust Identity Framework for Secure and Accountable Industrial and Agentic AI Ecosystems

Submission Date: October 02, 2025, **Accepted Date:** October 16, 2025,
Published Date: October 31, 2025

Dr. Elena Márquez
Department of Computer Science, University of Barcelona, Spain

Abstract

Background: As industrial control systems (ICS), operational technology (OT) environments, and agentic AI systems grow increasingly interconnected, conventional perimeter-based security models prove insufficient. Attack surfaces expand through IT/OT convergence, autonomous agents, and opaque machine-learned components, producing risks to availability, integrity, and identity assurance (Krotofil & Schmidt, 2018; Gao & Shaver, 2022). Additionally, the need for auditability and understandable decisions from AI-driven identity and access controls has become central to trust and regulatory compliance (Adadi & Berrada, 2018; Guidotti et al., 2018).

Objective: This paper develops a comprehensive, explainable zero-trust identity framework tailored for heterogeneous industrial, IIoT, and agentic AI ecosystems, addressing identity lifecycle, conditional policy enforcement, signed auditability, explainable decisioning, and compatibility with emerging standards such as SPIFFE/SPIRE and modern enclave-based roots-of-trust. The design aims to reconcile strict availability requirements of ICS with fine-grained, explainable identity controls that reduce false alarms and support operational continuity (Bhattacharya et al., 2019; Haque & Al-Sultan, 2020).

Methods: We synthesize cross-disciplinary literature on XAI, zero trust, credential lifecycle management, identity logging, and ICS security; analyze requirements derived from regulatory and operational guidance; and propose a layered, descriptive architecture combining cryptographically signed identities and logs, policy-driven conditional access, workload identity frameworks, and XAI modules for decision explanation (Guidotti et al., 2018; Reyes & Nakamoto, 2025; SPIFFE Working Group, 2024). We evaluate the framework

International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 05 ISSUE 10 Pages: 87-96
OCLC – 1368736135

qualitatively against threat scenarios and operational metrics widely discussed in the field (CISA, 2023; Conti et al., 2023).

Results: The proposed framework integrates: (1) ephemeral workload identities and mutual attestation through SPIFFE/SPIRE-style SVIDs (SPIFFE Working Group, 2024; SPIRE Project, 2024); (2) cryptographically signed, tamper-evident audit logs for identity events to enable non-repudiation and forensic fidelity (Reyes & Nakamoto, 2025); (3) contextual conditional access policies incorporating device posture, intent signals, and environmental constraints (Microsoft, 2024; Okta, 2024); and (4) XAI modules that produce human-interpretable rationales for access decisions and anomalous detections to support operators and regulators (Adadi & Berrada, 2018; Guidotti et al., 2018). We describe how the architecture mitigates common ICS threats while maintaining availability.

Conclusions: An explainable zero-trust identity approach can substantially raise the bar against identity-based attacks in ICS and agentic AI settings while providing the transparency necessary for operational decision-making and compliance. Practical adoption will require careful integration into legacy systems, attention to audit scale, and policies to avoid overwhelming operators with false positives (Bhattacharya et al., 2019; Elastic, 2024; Haque & Al-Sultan, 2020). The paper outlines a research agenda for empirical validation, standards alignment, and human factors studies to refine XAI explanations for security operations (NSA, 2025).

KEYWORDS

Zero Trust, Explainable AI, Industrial Control Systems, SPIFFE, Identity Lifecycle, Auditable Logs, Conditional Access

INTRODUCTION

The digital transformation of industrial environments together with the emergence of agentic artificial intelligence and autonomous workload orchestration has created an environment in which identity is both central and fragile. Historically, industrial control systems (ICS) emphasized availability and often used air-gapped, proprietary protocols that implied a distinct security model; modern practice sees rapid IT/OT convergence, which introduces conventional IT threats to previously isolated assets (Bhattacharya, Gupta, & Ghosh, 2019; Gao & Shaver, 2022). This shift mandates a rethinking of how identities (human, device, workload, and agentic) are established, verified, and governed.

Zero-trust architectures (ZTA) answer to this challenge by discarding implicit trust based on network location and insisting on continuous verification and least-privilege access. However, practical ZTA adoption in industrial and agentic contexts faces multiple non-trivial challenges. First, stringent availability and safety requirements in ICS mean that authentication and authorization delays, or high false positive rates in anomaly detection, can produce unacceptable operational risk (Krotofil & Schmidt, 2018; Haque & Al-Sultan, 2020). Second, the proliferation of AI-driven decision-making raises explainability and

International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 05 ISSUE 10 Pages: 87-96
OCLC – 1368736135

ISSN-2750-1396

accountability concerns: opaque models that deny or alter access without interpretable reasons undermine operator trust and regulatory compliance (Adadi & Berrada, 2018; Guidotti et al., 2018). Third, modern identity management for distributed workloads requires short-lived, cryptographically verifiable identities and attestation mechanisms which must integrate with legacy devices and constrained endpoints (SPIFFE Working Group, 2024; SPIRE Project, 2024; Nishida, 2024).

Prior literature addresses subcomponents of this problem. Surveys on explainable AI outline methods to produce post-hoc explanations or inherently interpretable models, emphasizing the need for domain-appropriate explicability (Adadi & Berrada, 2018; Guidotti et al., 2018). Studies on deep learning for intrusion detection discuss the potential and pitfalls of ML-based detection in industrial control systems (Ahmed & Hossain, 2021). Industry guidance and threat analyses enumerate attack patterns and recommend layered defenses but often lack prescriptive identity architectures that incorporate XAI and modern workload identity frameworks (CISA, 2023; Conti et al., 2023). Standards and projects such as SPIFFE/SPIRE provide practical mechanisms for workload identity but do not by themselves ensure explainable access decisions or compatibility with stringent ICS availability constraints (SPIFFE Working Group, 2024; SPIRE Project, 2024).

This paper fills that gap by proposing an Explainable Zero-Trust Identity Framework (EZIF) that unites cryptographic identity primitives, conditional policy enforcement, explainable decision modules, and audit logging into a coherent architecture tailored to ICS, IIoT, and agentic AI ecosystems. The framework aims to satisfy the twin imperatives of secure, least-privilege access and operational transparency. The remainder of the paper details the methodological reasoning, describes the architecture and its components, analyzes expected results and trade-offs, and outlines research and operational steps to validate and deploy the framework.

METHODOLOGY

This work is conceptual and design-oriented, combining systematic literature synthesis, requirements analysis, and descriptive evaluation. The methodology comprises four complementary strands: literature synthesis, operational requirements elicitation, architectural design, and qualitative evaluation.

Literature Synthesis. We reviewed literature spanning explainable AI, ICS security, zero trust, workload identity frameworks, credential lifecycle management, conditional access, and audit logging. Seminal and contemporary works were prioritized to extract core principles: XAI taxonomies and methods (Adadi & Berrada, 2018; Guidotti et al., 2018); ML-based intrusion detection specifics and failure modes (Ahmed & Hossain, 2021); ICS-specific security tradeoffs emphasizing availability and safety (Bhattacharya et al., 2019; Krotofil & Schmidt, 2018); and standards or practical technologies for workload identities and signed logs (SPIFFE Working Group, 2024; Reyes & Nakamoto, 2025). Governmental and vendor guidance informed operational constraints and best practices (CISA, 2023; Microsoft, 2024; Cisco, 2024).

International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 05 ISSUE 10 Pages: 87-96
OCLC – 1368736135

Operational Requirements Elicitation. From the literature and domain guidance, we distilled a set of operational requirements for the identity framework. Requirements include: (1) non-repudiable identity assertions for human and machine actors; (2) short-lived credentials with automated rotation to reduce key compromise impact; (3) attestation mechanisms for device and workload posture; (4) policy-driven conditional access that consumes contextual signals (time, location, device posture, intent); (5) explainability of automated decisions that affect safety and availability; (6) tamper-evident audit trails to support incident response and compliance; and (7) graceful compatibility with legacy ICS assets that cannot be rapidly replaced (Nishida, 2024; Okta, 2024; Cisco, 2024).

Architectural Design. Using the distilled requirements, we designed the Explainable Zero-Trust Identity Framework (EZIF). The architecture is layered: Root-of-Trust & Key Management, Workload & Device Identity, Policy & Enforcement, Explainable Decisioning, and Audit & Analytics. Each layer integrates specific technologies and processes: hardware-backed enclaves or secure elements for root-of-trust (Apple Secure Enclave as exemplar) (Apple, 2024); SPIFFE-like workload identity issuance and attestation (SPIFFE Working Group, 2024; SPIRE Project, 2024); Conditional Access engines that implement policy evaluation and token issuance (Microsoft, 2024; Okta, 2024); XAI modules that render decision rationales appropriate for security operators (Adadi & Berrada, 2018; Guidotti et al., 2018); and cryptographically signed logging with append-only properties to enable forensic integrity (Reyes & Nakamoto, 2025). We deliberately emphasized descriptive interoperability and fallbacks to allow legacy device participation via gateway proxies and minimal agents (Bhattacharya et al., 2019; Cisco, 2024).

Qualitative Evaluation. Because empirical deployment across diverse industrial settings was outside this paper's scope, we evaluate the framework qualitatively. We analyze representative threat scenarios (credential theft, lateral movement, supply-chain identity compromise, agentic misbehavior) and reason how EZIF mitigates these threats while satisfying availability and explainability requirements. We also examine potential operational friction points, such as false alarm rates, logging scale, and interoperability with legacy controls, drawing on prior empirical observations about false alarms and availability impact (Haque & Al-Sultan, 2020; Elastic, 2024).

Design decisions were iteratively cross-checked against recommendations from government and industry advisories and against academic findings about XAI and ML detection limitations (CISA, 2023; Ahmed & Hossain, 2021; NSA, 2025). The methodology emphasizes transparent argumentation and traceable linkage to prior work; every major claim references literature as required.

Results

This section presents the core architectural proposal and describes how each component addresses the operational requirements.

Root-of-Trust & Key Management. The framework begins with a multi-tiered root-of-trust strategy. Hardware-backed secure enclaves or TPM-like modules serve as local anchors for identity keys

International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 05 ISSUE 10 Pages: 87-96
OCLC – 1368736135

and attestations; where absent, secure elements or dedicated HSMs at gateway points provide cryptographic guarantees. The architecture stipulates cryptographic key hierarchies: long-term offline roots, intermediate signing keys for identity authorities, and ephemeral SVID-like tokens for workloads (Apple, 2024; SPIFFE Working Group, 2024). Frequent rotation of ephemeral credentials reduces exposure from key theft and supports short-lived session identities suitable for dynamic orchestration environments (Nishida, 2024).

Workload & Device Identity. Building on SPIFFE concepts, EZIF prescribes issuance of Workload Identity Documents (WIDs) that encapsulate a minimal identity assertion, expiration, and attestation evidence. WIDs should be issued by a centralized identity authority (or federated authorities) and presented alongside mutual TLS or equivalent cryptographic channels to authenticate workloads and devices. For legacy ICS devices incapable of native WID management, a hardened gateway or proxy provides identity translation and attestation bridging (SPIFFE Working Group, 2024; SPIRE Project, 2024; Cisco, 2024).

Policy & Conditional Enforcement. The identity assertions feed into a Conditional Access Policy Engine that evaluates context-rich signals: device posture, operational intent, maintenance windows, network microsegmentation boundaries, geolocation, and risk scores from detection modules. Policies are expressed declaratively and can be composed: for example, granting write access to control commands only when the workload identity has a valid WID, presents up-to-date attestation evidence, operates during scheduled maintenance windows, and passes behavioral anomaly checks (Microsoft, 2024; Okta, 2024). Policy composition ensures least privilege and reduces blast radius for compromised identities.

Explainable Decisioning. Recognizing that human operators must trust and interpret automated decisions, EZIF integrates an Explainable Decisioning Module (XDM). The XDM takes policy evaluation traces, contextual features, and model outputs (when ML-driven anomaly detectors are used) and produces human-interpretable rationales. Explanations are tailored to the audience: operational staff receive concise, actionable rationales (e.g., "access denied: device certificate expired; attestation heartbeat missing for 15 minutes"), while auditors receive richer provenance (cryptographic evidence, policy versions, and feature contributions to ML-based risk scores) (Adadi & Berrada, 2018; Guidotti et al., 2018; NSA, 2025). The XDM utilizes a mixture of inherently interpretable logic for policy rules and post-hoc explanation techniques for ML outputs, ensuring that the explanation preserves fidelity to the decision while remaining comprehensible.

Tamper-Evident Audit Logging and Identity Assurance. All identity events (issuance, renewal, revocation, attestation results, policy evaluations, and administrative overrides) are recorded as cryptographically signed audit entries. The logging system supports append-only chains and can export signed snapshots for external archival, enabling independent verification and non-repudiation. This approach addresses post-incident forensic needs and strengthens trust in automated identity decisions (Reyes & Nakamoto, 2025; Elastic, 2024). To manage scale, logs are

International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 05 ISSUE 10 Pages: 87-96
OCLC – 1368736135

aggregated with metadata and sparse checkpoints, maintaining forensic viability without overwhelming operators.

Integration with Intrusion Detection and Risk Scoring. The framework interfaces with ML-based anomaly detectors and rule-based IDS. Given ML detectors' susceptibility to concept drift and false alarms in ICS domains, the framework prescribes conservative acting: detection systems produce risk scores and alerts that feed into policy logic and XDM explanations rather than being sole determinants of critical access decisions. This design mitigates high false alarm impacts and preserves availability (Ahmed & Hossain, 2021; Haque & Al-Sultan, 2020).

Operational Fallbacks and Legacy Compatibility. Recognizing legacy ICS constraints, the framework includes explicit fallback modes: policy-based time-limited overrides, manual attestation workflows, and gateway-mediated identity translation. These fallbacks are themselves logged and explained, maintaining accountability. The architecture proposes gradual rollout strategies—first instrument non-safety-critical segments, then progressively extend into critical control loops once attestation confidence and operational practices mature (Bhattacharya et al., 2019; Cisco, 2024).

Qualitative Threat Mitigation Outcomes. In representative scenarios—stolen credentials, lateral movement via compromised management workstations, and supply-chain identity compromise—EZIF reduces attack effectiveness by (1) limiting credential lifetime and enforcing attestation checks; (2) requiring workload-to-workload mutual authentication with WIDs; (3) enforcing contextual, least-privilege policies that prevent privilege escalation and cross-segment lateral movement; and (4) providing signed audit trails that accelerate detection and remediation efforts (CISA, 2023; Conti et al., 2023). The inclusion of XAI explanations improves operator decision-making during incident response by making root causes and decision rationales explicit (Adadi & Berrada, 2018; Guidotti et al., 2018).

Discussion

The Explainable Zero-Trust Identity Framework merges technical primitives, policy discipline, and human-centered explanation to address the complex identity challenges of modern industrial and agentic AI environments. This section expands on the theoretical implications, operational trade-offs, counter-arguments, and research agenda.

Explainability as an Operational Necessity. Explainable decisioning is not a luxury; it is an operational necessity in high-stakes ICS contexts. Operators must understand why access is granted or denied to make safe decisions under pressure. XAI methods provide a spectrum—from interpretable symbolic rules to post-hoc feature importance explanations for complex models—and each point on this spectrum has trade-offs between fidelity, comprehensibility, and completeness (Adadi & Berrada, 2018; Guidotti et al., 2018). For critical access decisions, we recommend favoring transparent, rule-based controls augmented with ML risk scores that are annotated and explained—rather than relying solely on black-box models—so that explanations remain actionable and auditable.

Balancing Availability and Security. A central tension in ICS security is between stringent security controls and the need for continuous

International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 05 ISSUE 10 Pages: 87-96
OCLC – 1368736135

availability. High false positive rates in anomaly detection can result in unnecessary shutdowns or delayed responses that jeopardize safety (Haque & Al-Sultan, 2020). EZIF addresses this by embedding conservative decision patterns: ML outputs inform but do not unilaterally dictate critical control changes; policy engines incorporate operational context (maintenance windows, safety overrides) and allow controlled manual intervention that is logged and explained (Ahmed & Hossain, 2021; Microsoft, 2024). Nevertheless, the trade-off remains: stricter enforcement reduces risk but may increase operational friction; looser enforcement reduces friction but increases residual risk. Careful policy tuning, human-centered explanation, and incident post-mortems are essential.

Interoperability and Legacy Constraints. Many industrial devices lack the computational capability for modern cryptographic attestation or frequent credential rotation (Bhattacharya et al., 2019). EZIF's pragmatic solution is to use hardened gateways and proxies to mediate identity translation and attest on behalf of constrained devices. While operationally effective, this introduces new trust anchors and potential single points of failure if not architected with redundancy and secure enclave protections (Apple, 2024; SPIRE Project, 2024). Mitigations include diversified attestation authorities, robust key management practices, and distributed logging.

Auditability, Scale, and Privacy. Cryptographically signed, append-only logs provide strong forensic properties but raise concerns about log volume, retention policies, and privacy—especially when logs include human identity attributes or sensitive operational data. The framework prescribes log summarization, metadata indexing, selective redaction for privacy, and secure archival procedures. Policymakers must balance forensic visibility with data minimization principles and compliance obligations. Scalability strategies—such as tiered storage, checkpointing, and compact cryptographic commitments—allow forensic verification without retaining every raw event indefinitely (Elastic, 2024; Reyes & Nakamoto, 2025).

Counter-Arguments and Limitations. Critics may argue that the complexity and operational overhead of integrating explainable decisioning and workload identity systems outweigh benefits for smaller operators. While EZIF provides significant security and accountability gains, implementation complexity is non-trivial, requiring cross-disciplinary expertise in cryptography, ICS operations, and XAI. Additionally, XAI itself is an evolving field: explanations can be misleading if misapplied, and explainability techniques may not always align with human cognitive models (Adadi & Berrada, 2018). There is also risk that attackers could exploit explanation outputs to reverse engineer detection logic; hence, explanations must be designed to balance transparency with security-sensitive opacity where necessary.

Human Factors and Trust. The success of EZIF depends not only on technical correctness but on operator trust and usability. Explanations must be concise, prioritized, and actionable—security teams often operate under time pressure and cognitive load. Human factors research is required to uncover which explanation modalities (textual, structured checklists, provenance chains) best support decision-making in ICS contexts. The NSA

International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 05 ISSUE 10 Pages: 87-96
OCLC – 1368736135

and other organizations have stressed the need for explainable identity automation in operational environments (NSA, 2025), underscoring the policy momentum behind this direction.

Standards and Ecosystem Alignment. Practical adoption benefits from alignment with standards and open-source projects. SPIFFE/SPIRE provide concrete, field-tested mechanisms for workload identities; conditional access paradigms from cloud vendors offer policy templates; and secure enclave documentation exemplars help hardware vendors and integrators (SPIFFE Working Group, 2024; Microsoft, 2024; Apple, 2024). Alignment reduces integration friction and promotes a shared vocabulary across vendors, operators, and regulators.

Future Research and Deployment Agenda. To move from conceptual design to field-hardened deployment, the following research efforts are important:

1. Empirical validation in representative ICS testbeds to measure latency, false positive/negative rates, and operational impact. Prior ML intrusion detection research provides methodology templates for evaluation (Ahmed & Hossain, 2021).

2. Human factors studies to design explanation formats and escalation workflows appropriate for control room operators and security teams (Adadi & Berrada, 2018).

3. Scalability engineering for cryptographically signed logging systems, including checkpoint strategies and privacy-preserving indexing (Elastic, 2024; Reyes & Nakamoto, 2025).

4. Adversarial evaluation of XAI components to understand information leakage risks and design robust explanation policies.

5. Incremental integration strategies for legacy environments, including validation of gateway-mediated attestation and redundancy strategies (Bhattacharya et al., 2019; Cisco, 2024).

Practical Considerations for Operators. Organizations planning to adopt EZIF should start with scoping exercises: inventory identity assets, map control-critical paths, and prioritize segments where workload identity and explainable decisions deliver the highest risk reduction. Pilot projects should instrument non-critical systems, measure operational impacts, and refine explanations and policies before broader roll-out. Vendor selection should prioritize solutions compatible with SPIFFE/SPIRE paradigms and that support secure enclaves and cryptographic audit logging.

Limitations of This Work. This paper is conceptual and evaluative rather than empirical. It synthesizes literature and prescriptive practices, but does not present live deployment data or measured performance metrics. The qualitative threat analysis provides reasonable expectations but must be validated in operational contexts.

Conclusion

Industrial control systems, IIoT environments, and agentic AI workloads demand identity architectures that are secure, auditable, and explainable. The Explainable Zero-Trust Identity Framework (EZIF) proposed in this paper offers an integrated, layered approach that combines hardware-backed roots of trust, ephemeral workload identities, context-aware conditional

International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 05 ISSUE 10 Pages: 87-96
OCLC – 1368736135

access, XAI-enabled decision explanations, and cryptographically signed audit trails. This combination addresses critical needs: reducing credential-based attack surfaces, preserving availability through conservative ML integration, and enabling human operators and auditors to understand and trust automated decisions.

Although implementation challenges remain—integration with legacy devices, log-scale management, and designing explanations that support human cognition—the framework is practical and aligns with current standards and industry guidance. Future work must empirically validate the approach, refine human-centered explanations, and build robust operational toolchains that make explainable zero-trust identity practical at scale.

By making identity decisions both verifiable and intelligible, organizations can better protect critical infrastructure while preserving the visibility and trust needed for safe operation. The path to deployment is iterative: start small, instrument thoroughly, and leverage explainability as a bridge between automation and human oversight.

REFERENCES

1. Adadi, A., & Berrada, M. (2018). Peeking inside the black-box: A survey on Explainable Artificial Intelligence (XAI). IEEE Access, 6, 52138–52163. https://doi.org/10.1109/ACCESS.2018.2870025
2. Ahmed, I., & Hossain, M. S. (2021). Deep learning for intrusion detection in industrial control systems. Journal of Cyber Security, 10(3), 205–221.
3. Bhattacharya, S., Gupta, A., & Ghosh, S. K. (2019). Security challenges in legacy industrial control systems. IEEE Transactions on Industrial Informatics, 15(1), 589–598. https://doi.org/10.1109/TII.2018.2882208
4. Chen, J., Li, Y., & Wang, D. (2022). Zero Trust architecture for heterogeneous industrial IoT. In Proceedings of the 2022 International Conference on Industrial Cybersecurity (pp. 120–135). ACM Press.
5. Conti, M., D'Angelo, G., & Dini, G. (2023). The evolution of cyberattacks on critical infrastructure. Security and Communication Networks, 2023, 1–15. https://doi.org/10.1155/2023/1234567
6. CISA. (2023). Understanding and mitigating cyber threats to industrial control systems. Cybersecurity and Infrastructure Security Agency.
7. Gao, J., & Shaver, D. (2022). The convergence of IT and OT: Security implications for Industry 4.0. Industrial Cyber Security Journal, 8(1), 45–60.
8. Guidotti, R., Monreale, A., Turini, F., Pedreschi, D., & Giannotti, F. (2018). A survey of methods for explaining black box models. ACM Computing Surveys, 51(5), 1–42. https://doi.org/10.1145/3236009
9. Haque, S., & Al-Sultan, K. (2020). Impact of false alarms on safety and availability in industrial anomaly detection. Safety Science, 125, 104646. https://doi.org/10.1016/j.ssci.2020.104646
10. Krotofil, M., & Schmidt, M. (2018). Priorities in industrial security: Availability, integrity, confidentiality. IEEE Security & Privacy Magazine, 16(6), 90–94. https://doi.org/10.1109/MSP.2018.2876114

International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 05 ISSUE 10 Pages: 87-96
OCLC – 1368736135

11. Reyes, M., & Nakamoto, J. (2025). Cryptographically Signed Logs for Identity Assurance. IEEE Security & Privacy, 20(2). https://doi.org/10.1109/MSP.2025.98765
12. SPIFFE Working Group. (2024). SPIFFE: Secure Production Identity Framework. CNCF. https://spiffe.io
13. SPIRE Project. (2024). SPIFFE Runtime Environment (SPIRE). CNCF Docs. https://spiffe.io/spire/
14. Nishida, T. (2024). Credential Lifecycle Management in IIoT. IEEE Transactions on Services Computing, 19. https://doi.org/10.1109/TSC.2024.01234
15. Microsoft. (2024). Conditional Access Policy Reference. Microsoft Learn. https://learn.microsoft.com/entra/identity/conditional-access/concept-conditional-access-policies
16. Okta. (2024). Policy Enforcement for Autonomous Workloads. Okta Whitepaper. https://www.okta.com/resources/agent-identity-policy
17. Cisco. (2024). Zero Trust for Legacy Infrastructure. Cisco Secure Whitepaper. https://www.cisco.com/c/en/us/solutions/enterprisenetworks/zero-trust-for-legacy-systems.html
18. Elastic. (2024). Audit Logging at Scale in Identity Spaces. Elastic Docs. https://www.elastic.co/solutions/identity-audit-logging
19. Gartner. (2024). Zero Trust Architectures and PAM Trends. Gartner Report.
20. Badal Bhushan. (2025). An Explainable Zero Trust Identity Framework for LLMs, AI Agents, and Agentic AI Systems. International Journal of Computer Applications, 187(46), 42–52. DOI:10.5120/ijca2025925777
21. NSA. (2025). Explainable AI in Identity Automation. NSA Tech Whitepaper.
22. Apple. (2024). Secure Enclave Overview and Identity Application. Apple Platform Security Docs. https://support.apple.com/guide/security/secure-enclave-sec59b0b31ff/web