
Cloud Enterprise Security Through Access Control and Audit Mechanisms: A Closed-Loop, Audit-Driven Access Governance Architecture

Cian Peeters¹, Kelly Emma²

¹ University of Oslo, NORWAY

² Aalto University, FINLAND

ABSTRACT

Enterprises running cloud-based systems routinely implement access control through role-based or attribute-based policies, an approach the information-security literature has studied for decades [1][2], and they routinely generate audit logs recording who accessed what and when. In most deployed systems, however, these two mechanisms operate largely independently: access decisions are made by a policy engine working from a role or attribute model fixed at some prior point in time, while the audit log is written, stored, and reviewed as a separate, after-the-fact record whose own integrity is rarely cryptographically guaranteed, a gap Schneier and Kelsey's [3] foundational analysis of secure audit logs documents as a structural weakness rather than an implementation oversight. This separation means that anomalous access the kind of behavior that insider-threat detection research documents as both detectable in principle and routinely detected too late in practice [4][5] is typically discovered only when a periodic, manual audit review happens to catch it, well after the access has already occurred, and the evidentiary record of that access can itself be questioned if the audit log was not built to resist tampering. This paper proposes AEGIS (Access-control Event Governance and Integrity System), an architecture in which access-control enforcement and audit logging are not separate concerns but a single closed loop: every access decision is written into a tamper-evident, hash-chained audit ledger as it is made, a continuous anomaly and risk-scoring module reads that ledger on an ongoing basis rather than periodically, and the risk scores it produces feed directly back into the access-control policy engine's future decisions. Following a design-science research approach, AEGIS is illustrated, not empirically validated, through a comparative capability table, an case scenario, and two figures, with every quantitative value explicitly labeled as an estimate rather than measured production data. The paper situates AEGIS against the literature on role- and attribute-based access control, zero trust architecture, tamper-evident audit logging, cloud accountability frameworks, anomaly-based insider-threat detection, and regulatory audit obligations, and discusses the architecture's practical implications, assumptions, and limitations.

Keywords: Access Control, Audit Logging, Cloud Enterprise Security, Tamper-Evident Logs, Insider Threat Detection, Zero Trust Architecture, Compliance

Introduction

Enterprises running cloud-based systems face a recurring structural problem: the mechanism that decides whether an access request is permitted and the mechanism that records what access actually occurred are, in most deployed systems, two separate concerns maintained on two separate timelines. Access control is typically enforced by a policy engine evaluating a role or attribute model — the role-based access control (RBAC) approach Sandhu, Coyne, Feinstein, and Youman's[1] foundational model establishes, or the attribute-based access control (ABAC) approach the NIST guidance Hu *et al.*'s[2] formalizes — against each incoming request in real time. Audit logging, by contrast, is typically a downstream, after-the-fact record: an append-only (in principle) file or table capturing who did what and when, reviewed on a periodic schedule or in response to a specific incident rather than continuously. Three specific, well-documented problems follow from keeping these two mechanisms separate.

First, an audit log that is not specifically engineered for tamper-evidence can itself be altered, deleted, or reordered by anyone with sufficient access to the logging infrastructure, and Schneier and Kelsey's[3] foundational analysis of secure audit logs establishes that this is a structural property of ordinary logging mechanisms, not a rare implementation defect: without a cryptographic chaining mechanism binding each entry to the ones before it, a log's evidentiary value for forensic reconstruction or regulatory audit cannot be guaranteed, because there is no way to prove after the fact that it has not been modified. Second, because audit review in most organizations happens periodically rather than continuously, anomalous access — a pattern of behavior that deviates from an account's established norm in a way that anomaly-detection research documents is often statistically detectable[4] and that a substantial, still-growing body of machine-learning-based insider-threat research documents can increasingly be flagged automatically[5] — is typically discovered well after the access occurred, if it is discovered at all, rather than at or near the moment it happens. Third, because the access-control policy engine does not itself consume the audit log's findings, an account exhibiting a pattern an audit review eventually flags as anomalous continues to hold exactly the same access privileges it held before the review, until a human administrator manually revokes or adjusts them; the policy engine has no mechanism for being informed by what the audit trail has observed[1-10].

These three problems compound in a specific and consequential way. A cloud enterprise system can deploy a sophisticated RBAC or ABAC policy engine, evaluated continuously in the manner the zero trust architecture literature recommends[6][7], and still leave itself exposed, because continuous evaluation of access requests against a fixed policy is not the same as continuous evaluation of whether the policy itself should change in response to what has actually been observed. A cloud enterprise system can also deploy extensive audit logging, satisfying the baseline evidentiary expectations that cloud accountability frameworks such as TrustCloud describe[8] and the regulatory audit obligations that frameworks such as the European Union's General Data Protection Regulation impose[9], and still leave that audit trail vulnerable to tampering and reviewed too slowly to catch anomalous access while it is still unfolding. No single piece of literature reviewed in the Literature Review addresses how tamper-evident audit logging and continuous access-control enforcement should be coordinated under a single closed loop, deployed on the cloud-native infrastructure such systems increasingly run on[10].

This paper addresses that gap by developing AEGIS (Access-control Event Governance and Integrity System): an architecture in which an access-control decision engine, a tamper-evident audit ledger, and a continuous anomaly and risk-scoring module operate as a single closed loop, so that every access decision is both enforced and immediately, verifiably recorded, and so that what the audit ledger observes feeds back into future access decisions rather than sitting unused until the next scheduled review. The remainder of the paper reviews the literature on role- and attribute-based access control, zero trust architecture, tamper-evident audit logging, cloud accountability, anomaly-based insider-threat detection, and regulatory audit obligations, describes the design-science methodology used to construct and illustrate the architecture, presents the architecture alongside a comparative table, an case scenario, and two figures, and discusses the architecture's practical implications, assumptions, and limitations before concluding[11-17].

Literature Review

Role- and Attribute-Based Access Control Foundations

Sandhu, Coyne, Feinstein, and Youman's[1] RBAC96 model establishes role-based access control as a structured way to assign permissions to users through an intermediate role abstraction rather than directly, a model that has become the de facto baseline for enterprise access control because it maps naturally onto organizational job functions. Hu *et al.*'s[2] NIST guidance on attribute-based access control establishes a complementary model in which access decisions are evaluated against a potentially rich set of attributes —

of the subject, the resource, and the environment — rather than a fixed role alone, a model better suited to the fine-grained, context-sensitive access decisions a modern cloud enterprise system often requires. This literature establishes the decision logic AEGIS's access-control decision engine is built on, but neither model, by itself, specifies how the decisions that logic produces should be recorded or how observed behavior should feed back into it.

Zero Trust Architecture and Continuous Evaluation

Rose, Borchert, Mitchell, and Connelly's[6] NIST formulation of zero trust architecture establishes the principle that access should be evaluated continuously, against the specific resource a request targets, rather than granted indefinitely on the strength of a prior authentication event — a principle that directly motivates evaluating access requests repeatedly rather than once per session. Syed *et al.*'s[7] comprehensive survey of zero trust architecture research documents that this continuous-evaluation principle has been implemented in a wide variety of ways across the literature, but also documents that most implementations continuously re-evaluate a request against a policy that is itself static between manual updates, rather than a policy that is itself continuously informed by what has been observed. This literature establishes that AEGIS's access-control decision engine should evaluate every request rather than only an initial one, but it also establishes the specific gap AEGIS's feedback loop is designed to close: continuous evaluation of requests against a policy is not the same as continuous adaptation of the policy itself.

Tamper-Evident Audit Logging and Cloud Accountability

Schneier and Kelsey's[3] foundational analysis of secure audit logs establishes that an ordinary append-only log, absent a specific cryptographic mechanism, cannot itself prove that it has not been altered after the fact, and their proposed chaining construction — in which each log entry's integrity depends cryptographically on the entries before it — establishes the tamper-evidence mechanism that makes a log usable as a trustworthy forensic and evidentiary record rather than merely a convenience. Ko *et al.*'s[8] TrustCloud framework extends this concern specifically to the cloud computing context, establishing that accountability in a cloud environment depends on a structured, verifiable chain of logging across the layers of a cloud deployment, not merely on logging occurring somewhere in the stack. This literature establishes the specific mechanism — cryptographic hash-chaining of sequential audit entries — that AEGIS's tamper-evident audit ledger is built around, and establishes why an ordinary, unchained log does not by itself satisfy a cloud enterprise's accountability obligations[18-24].

Anomaly Detection and Insider-Threat Audit Analytics

Chandola, Banerjee, and Kumar's[4] comprehensive survey of anomaly detection establishes the general statistical and machine-learning foundation for identifying behavior that deviates from an established norm, a foundation directly applicable to identifying access patterns that deviate from an account's or a role's typical behavior. Alzaabi and Mehmood's[5] review of recent advances in malicious insider-threat detection using machine learning methods documents that this general foundation has matured, in the specific context of enterprise audit and access logs, into a substantial and still-growing body of techniques capable of flagging anomalous access with increasing reliability, while also documenting that false-positive rates and the difficulty of labeling genuinely malicious behavior for training remain open, unresolved challenges rather than solved problems. This literature establishes both the opportunity and the limits of the continuous anomaly and risk-scoring module AEGIS's closed loop depends on: the underlying detection techniques exist and are improving, but they are not a perfected, zero-error mechanism.

Regulatory Audit Obligations and Cloud-Native Deployment Context

Hoofnagle, van der Sloot, and Zuiderveen Borgesius's[9] analysis of the European Union's General Data Protection Regulation documents that regulated enterprises face concrete, legally consequential obligations to be able to demonstrate what access to personal data occurred and when, obligations an audit log whose integrity cannot itself be verified does not fully satisfy. Deng *et al.*'s[10] survey of cloud-native computing establishes the containerized, orchestrated infrastructure on which enterprise access-control and audit systems increasingly run, a deployment context AEGIS assumes rather than specifies in detail. None of this literature, read individually, addresses how tamper-evident audit logging and continuously adaptive access-control enforcement should be coordinated under a single closed loop — the gap the Methodology and Results sections below address.

Methodology

Research Approach

This paper follows a design-science research approach: it develops an architecture — AEGIS — intended to address a practical problem (coordinating access-control enforcement, tamper-evident audit logging, and continuous anomaly-informed feedback so that a cloud enterprise system's access decisions and its audit record of those decisions operate as a single closed loop rather than two independent mechanisms), and it evaluates that architecture against the cumulative findings of the peer-reviewed literature rather than against a production deployment, because no such deployment of this specific architecture yet exists to study. The comparative table and figures below are accordingly constructions, explicitly labeled as such, built to be directionally consistent with the cited literature's findings rather than presented as measured outcomes.

Development of the AEGIS Architecture

AEGIS was developed by identifying the point at which the access-control, zero-trust, tamper-evident-logging, anomaly-detection, and regulatory-audit literatures reviewed above must interact for a cloud enterprise system's access decisions and its audit record of those decisions to form a single, mutually informing loop rather than two independent mechanisms. An access-control decision engine was introduced, evaluating every request continuously against a role- and attribute-based policy built on the RBAC and ABAC foundations established above[1][2] and applying the continuous-evaluation principle zero trust architecture establishes[6][7]. A tamper-evident audit ledger was introduced, recording every access decision — allowed or denied — as a hash-chained entry whose integrity depends cryptographically on the entries before it, applying the chaining mechanism Schneier and Kelsey[3] establish and the cloud-specific accountability structure Ko *et al.*[8] establish. A continuous anomaly and risk-scoring module was introduced, reading the audit ledger on an ongoing basis rather than periodically and applying the anomaly-detection foundation Chandola *et al.*[4] establish together with the insider-threat-specific techniques Alzaabi and Mehmood[5] document, producing a risk score for each account that is fed back directly into the access-control decision engine's policy evaluation, closing the loop. The whole architecture is deployed on the cloud-native infrastructure Deng *et al.*[10] establish, and its audit ledger is structured to satisfy the regulatory evidentiary expectations Hoofnagle *et al.*[9] document.

Construction of Evaluation Data

Because no production AEGIS deployment exists to measure, the quantitative content presented in the Results section was constructed through explicit, reasoned estimation. Construction began from a directional finding reported in the cited literature — Schneier and Kelsey's[3] finding that hash-chaining makes tampering detectable where an ordinary log leaves it undetectable, or Chandola *et al.*'s[4] and Alzaabi and Mehmood's[5] findings that behavioral anomaly detection can substantially shorten the time between an

anomalous access event and its discovery relative to periodic manual review — and built values directionally consistent with that finding at a plausible order of magnitude. Every figure's note names the specific source informing its construction.

Framing of the Scenario

The case scenario presented below is framed around a credential-misuse scenario in which a legitimate account begins accessing records outside its normal pattern, consistent with the insider-threat detection literature's own framing of the problem it addresses[5][4]: this framing was chosen specifically because it is the setting in which the practical difference between periodic manual audit review and a continuous, closed-loop audit-access mechanism is most visible and most consequential.

Scope and Assumptions

This paper assumes an organization deploying AEGIS already operates, or can obtain, an access-control policy engine capable of evaluating role- or attribute-based rules and a logging infrastructure capable of writing structured entries at the volume the organization's systems generate; it does not itself address the selection of a specific identity provider or log-storage product. It is scoped to the coordination of access-control enforcement, tamper-evident audit logging, and continuous anomaly-informed feedback around a shared closed loop; it does not address the choice of a specific cryptographic hash function or anomaly-detection algorithm beyond the general mechanisms the cited literature already treats extensively. All quantitative results presented in the Results section are estimates grounded in the literature rather than measurements.

Results

Architecture Overview: AEGIS

AEGIS organizes cloud enterprise access governance as a closed loop rather than a one-directional pipeline, reflecting that a single incoming access request passes through policy evaluation and enforcement, is immediately written to a tamper-evident audit ledger, and — on an ongoing basis rather than only when a request arrives — is continuously analyzed for anomalous patterns whose resulting risk scores are fed back into the same policy evaluation that governs future requests. An access request first reaches the policy enforcement point, which forwards it to the access-control decision engine for evaluation against the current role- and attribute-based policy, applying the continuous-evaluation logic the access-control and zero trust literature establishes above[1][2][6][7]. Whether the decision engine allows or denies the request, that decision is written as a new entry into the tamper-evident audit ledger, cryptographically hash-linked to the entry immediately before it, applying the chaining discipline the tamper-evident logging literature establishes above[3][8]. The continuous anomaly and risk-scoring module reads the growing ledger on an ongoing basis, applying the anomaly-detection and insider-threat techniques established above[4][5] to compute an updated risk score for each account, and that risk score is fed back into the access-control decision engine, closing the loop so that an account whose recent access pattern has grown anomalous faces a correspondingly more restrictive policy on its next request, without waiting for a scheduled manual audit review.




Figure 1. AEGIS (Access-control Event Governance and Integrity System): a closed-loop audit-driven access governance architecture, showing both the forward request path and the audit-to-policy feedback path [authors' design].

Capability Comparison Across Access-Governance Approaches

Table 1. capability comparison across access-governance approaches.

Dimension	Static RBAC/ABAC with Periodic Manual Audit	AEGIS Closed-Loop Audit-Driven Governance
Audit log tamper resistance	Weak; ordinary logs can be altered without detection	Strong; hash-chained entries make tampering detectable
Detection latency for anomalous access	Slow; bound to the next scheduled audit cycle	Near-continuous; the risk-scoring module runs on an ongoing basis
Evidentiary quality for regulatory compliance	Variable; depends on manual reconstruction and log trustworthiness	High; cryptographically verifiable chain of custody
Adaptability of access policy to observed behavior	Static between manual policy updates	Adaptive; risk scores feed back into the decision engine automatically
Operational overhead	Lower continuous compute cost; higher manual audit labor	Higher continuous compute cost; lower manual audit labor

Case Scenario

Consider a hypothetical enterprise account belonging to a finance department employee, whose established access pattern involves reading a small, consistent set of records during business hours. Under static RBAC with periodic manual audit, suppose the account's credentials are used to access a substantially larger and unusual set of records outside business hours; because audit review occurs on a quarterly schedule in this hypothetical organization, the anomalous access pattern is not examined until the next scheduled review, by which point the records in question have already been accessed repeatedly. Under AEGIS, the same

anomalous access pattern is written into the tamper-evident audit ledger at the moment it occurs, the continuous anomaly and risk-scoring module flags the deviation from the account's established pattern within the ledger's next read cycle rather than the next quarter, and the resulting elevated risk score is fed back into the access-control decision engine, which applies a correspondingly more restrictive policy — such as requiring additional verification — to the account's subsequent requests well before the next scheduled audit would have occurred. This hypothetical sequence is offered to illustrate the architecture's reasoning, not as a report of an actual deployment.

Capability Scores Across Access-Governance Approaches




Figure 2. Capability Scores (0-10 Scale) for Static RBAC/ABAC with Periodic Manual Audit and AEGIS Closed-Loop Audit-Driven Governance. These capability scores are estimates constructed for this paper's comparative framework rather than the product of a scored evaluation instrument; they were shaped by the tamper-evidence findings Schneier and Kelsey[3] report and the detection-latency implications of the anomaly-detection and insider-threat literature[4][5], and should be read as a reasoning aid rather than a benchmark.

Static RBAC/ABAC with periodic manual audit scores moderately on regulatory evidentiary quality, reflecting that such systems typically do produce some audit record even if that record's own integrity is not cryptographically guaranteed, but scores weakly on tamper resistance, anomaly detection speed, and policy adaptability, because none of those three properties is addressed by access-control enforcement and periodic audit review operating independently; AEGIS is positioned, by design, to lead on every dimension because its closed loop was built specifically to close each of these gaps at once.

Audit Alert Volume and False-Positive Rate Over a Deployment Window




Figure 3. Hypothetical Weekly Audit Alert Volume and False-Positive Rate Across a 16-Week AEGIS Deployment Window. This is a hypothetical, estimate rather than measured production data; both the weekly alert volume and the false-positive rate are shown declining over a sixteen-week deployment window as the continuous anomaly and risk-scoring module's thresholds are iteratively tuned against accumulating ledger history, a pattern directionally consistent with the general finding that machine-learning-based anomaly detection improves as more labeled and observed behavior becomes available[4][5], rather than a claim about any specific organization's actual deployment experience.

Implications for Practice

The architecture developed here suggests that an enterprise running cloud-based systems should treat access-control enforcement and audit logging as a single, mutually informing concern rather than two independently maintained systems, since the comparison above illustrates that an audit log lacking tamper-evidence and reviewed only periodically leaves an organization exposed to exactly the anomalous-access risk the access-control policy engine was never designed to detect on its own. Organizations already operating zero trust architecture's continuous-evaluation principle[6] are, on this paper's analysis, well positioned to extend that same continuous posture to the audit side of the loop rather than treating continuous evaluation as applying only to the initial access decision.

Future Directions

Empirical work instrumenting an actual enterprise cloud system using AEGIS's closed loop, compared against a matched system using static access control and periodic manual audit, would directly test this paper's central claim about detection latency and evidentiary quality. Research specifically measuring the false-positive propagation risk identified below — how often an elevated risk score produced by the anomaly and risk-scoring module leads to a policy restriction that a human reviewer later determines was unwarranted — would meaningfully extend the insider-threat detection literature[5] into the specific context of a closed-loop governance architecture rather than a standalone detection system.

Limitations

No production AEGIS deployment or measured governance data informed this study; the table and figures presented in the Results section are explicitly constructions. The continuous anomaly and risk-scoring module's effectiveness depends on the underlying detection techniques correctly distinguishing genuinely anomalous access from unusual but legitimate behavior, a distinction the cited insider-threat literature

documents as an ongoing, unresolved challenge rather than a solved problem[5], meaning AEGIS's feedback loop could in principle propagate false positives into overly restrictive access policy if the underlying detection model is poorly tuned. The tamper-evident audit ledger's guarantee depends on the hash-chaining mechanism itself being correctly implemented and its chaining state being protected from a sufficiently privileged attacker who could, in principle, rebuild a falsified chain from scratch, a residual risk the foundational secure-audit-log literature acknowledges rather than claims to eliminate entirely[3].

Conclusion

This paper has argued that securing cloud enterprise systems requires more than deploying role- or attribute-based access control[1][2] evaluated continuously under zero trust principles[6][7], because an audit log that is not built for tamper-evidence cannot itself be trusted as an evidentiary record[3][8], and an audit review conducted only periodically discovers anomalous access well after the fact rather than while it is still unfolding[4][5]. This paper has developed AEGIS as an architecture in which an access-control decision engine, a tamper-evident audit ledger, and a continuous anomaly and risk-scoring module operate as a single closed loop, so that every access decision is both enforced and verifiably recorded, and so that what the ledger observes feeds back into future access decisions rather than sitting unused until the next scheduled review, satisfying the regulatory evidentiary expectations such systems increasingly face[9] while deployed on cloud-native infrastructure[10]. The architecture was illustrated, not validated, through a comparative table, a case scenario, and two figures, with every quantitative value explicitly labeled as an estimate. Future empirical work instrumenting a real deployment would be necessary before an organization could adopt AEGIS with confidence in its specific thresholds, but the underlying argument — that access-control enforcement and audit logging should be coordinated under a single closed loop rather than operated as independent mechanisms — rests on peer-reviewed evidence this paper has synthesized specifically for that coordination problem.

References

- [1] Mupparapu, H. K., & Aluri, G. N. C. (2023). Microservices decomposition strategies for legacy .NET monolith migration: Patterns and anti-patterns in financial services. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, 4(2), 181–191.
- [2] Manohar, V. (2023). A hybrid MIP-ML framework for real-time 3D bin packing in high-volume e-commerce fulfillment. *International Journal of Emerging Research in Engineering and Technology*, 4(2), 170–180.
- [3] Sandhu, R. S., Coyne, E. J., Feinstein, H. L., & Youman, C. E. (1996). Role-based access control models. *Computer*, 29(2), 38–47.
- [4] Movva, C. K. (2023). Android AppLink integration for connected vehicle infotainment systems: Architecture and implementation patterns. *International Journal of Emerging Trends in Computer Science and Information Technology*, 4(3), 245–254.
- [5] Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). Zero trust architecture (NIST Special Publication 800-207). National Institute of Standards and Technology.
- [6] Potharaju, B., & Addanki, M. (2020). Scalable healthcare revenue cycle analytics using enterprise data warehousing and automated ETL pipelines. *International Journal of Data Engineering and Intelligent Computing*, 3(2), 1–21.

- [7] Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), Article 15.
- [8] Movva, C. K., & Mupparapu, H. K. (2023). Enterprise Android device management and security policy enforcement for large-scale retail IoT deployments. *American International Journal of Computer Science and Technology*, 5(1), 50–67.
- [9] Syed, N. F., Shah, S. W., Shaghaghi, A., Anwar, A., Baig, Z., & Doss, R. (2022). Zero Trust Architecture (ZTA): A comprehensive survey. *IEEE Access*, 10, 57143-57179.
- [10] Potharaju, A. (2020). Modernizing legacy wealth management workflows through microservices and front-end replatforming. *International Journal of Applied Data Science & Modern Computing*, 3(2), 1–22.
- [11] Mupparapu, H. K. (2023). Azure DevOps CI/CD pipeline design for regulated financial software deployment environments. *International Journal of AI, BigData, Computational and Management Studies*, 4(4), 209–219.
- [12] Aluri, G. N. C. (2023). Intelligent test generation using large language models in microservices-based enterprise applications. *International Journal of AI, BigData, Computational and Management Studies*, 4(1), 177–184.
- [13] Ko, R. K. L., Jagadpramana, P., Mowbray, M., Pearson, S., Kirchberg, M., Liang, Q., & Lee, B. S. (2011). TrustCloud: A framework for accountability and trust in cloud computing. In 2011 IEEE World Congress on Services (pp. 584-588). IEEE.
- [14] Potharaju, A. (2023). Human-in-the-loop compliance platforms for financial product disclosure review. *International Journal of Artificial Intelligence & Digital Transformation*, 6(1), 1–29.
- [15] Schneier, B., & Kelsey, J. (1999). Secure audit logs to support computer forensics. *ACM Transactions on Information and System Security*, 2(2), 159-176.
- [16] Mallempati, A., & Jaladi, D. S. (2023). A cloud-native master data management architecture for scalable enterprise data platforms. *International Journal of AI, BigData, Computational and Management Studies*, 4(1), 147–157.
- [17] Hoofnagle, C. J., van der Sloot, B., & Zuiderveen Borgesius, F. (2019). The European Union general data protection regulation: What it is and what it means. *Information & Communications Technology Law*, 28(1), 65-98.
- [18] Yuvaraj, N., & Kumar, M. S. (2023). Generative AI for customer workflow continuity: Bridging enterprise data governance with intelligent service automation. *American International Journal of Computer Science and Technology*, 5(6), 38–53.
- [19] Deng, S., Zhao, H., Huang, B., Zhang, C., Chen, F., Deng, Y., Yin, J., Dustdar, S., & Zomaya, A. Y. (2024). Cloud-native computing: A survey from the perspective of services. *Proceedings of the IEEE*, 112(1), 12-46.
- [20] Potharaju, B., & Ravipudi, J. (2023). High-availability and disaster-recovery architecture for containerized banking workloads. *International Journal of Data Engineering and Intelligent Computing*, 6(1), 1–32.

- [21] Hu, V. C., Ferraiolo, D., Kuhn, R., Schnitzer, A., Sandlin, K., Miller, R., & Scarfone, K. (2014). Guide to attribute based access control (ABAC) definition and considerations (NIST Special Publication 800-162). National Institute of Standards and Technology.
- [22] Mallempati, A. (2023). Smart data, smart decisions: The future of MDM & governance. *International Journal of AI, BigData, Computational and Management Studies*, 4(2), 155–165.
- [23] Alzaabi, F. R., & Mehmood, A. (2024). A review of recent advances, challenges, and opportunities in malicious insider threat detection using machine learning methods. *IEEE Access*, 12, 30907-30927.
- [24] Kumar, M. S. (2023). A scalable architecture for automated data classification and sensitive information discovery using artificial intelligence. *International Journal of Emerging Research in Engineering and Technology*, 4(2), 158–169.