Journal of Contemporary Science and Technology Management
Volume 02, Issue 04, December 2026

Forensic-Driven Cyberattack Attribution and Evidence Correlation Using Generative AI in Enterprise Cloud Environments

Dr. Dinesh Kumar Saini

Faculty of Computing and Information Technology, Sohar University, Oman

DOI: https://doi.org/10.66592/v2gh5x12

Abstract

The increasing sophistication of cyberattacks targeting enterprise cloud environments has created significant challenges for digital forensic investigation, evidence correlation, and reliable attack attribution. Distributed cloud infrastructure generates enormous volumes of security logs, authentication records, network traffic, endpoint activities, and application events, making manual investigation time-consuming and complex. This research proposes a forensic-driven cyberattack attribution and evidence correlation framework integrating Generative Artificial Intelligence (GenAI), machine learning, behavioral analytics, and knowledge graph technologies. The proposed framework collects digital evidence from heterogeneous enterprise cloud platforms and performs preprocessing, normalization, timestamp alignment, and forensic integrity verification. Generative AI analyzes correlated security events, reconstructs potential attack sequences, identifies relationships among suspicious entities, and generates contextual explanations to support forensic investigations. Machine learning techniques classify malicious activities, detect anomalous behaviors, and identify potential attack indicators, while knowledge graphs establish relationships among users, devices, workloads, network connections, and compromised resources. Evidence correlation mechanisms combine multiple forensic artifacts to improve investigative accuracy and reduce fragmented interpretations. Furthermore, explainable attribution mechanisms evaluate behavioral similarities, infrastructure relationships, and threat intelligence indicators without treating automated predictions as definitive identification of attackers. The framework incorporates evidence provenance, cryptographic hashing, access controls, and audit trails to preserve forensic reliability. This research aims to improve investigation efficiency, evidence traceability, attack reconstruction, and cybersecurity decision-making across distributed enterprise cloud environments.

Journal of Contemporary Science and Technology Management Volume 02, Issue 04, December 2026

Keywords: Digital Forensics, Cyberattack Attribution, Generative Artificial Intelligence, Evidence Correlation, Enterprise Cloud Security, Knowledge Graphs, Machine Learning, Threat Intelligence, Attack Reconstruction, Forensic Evidence Integrity.

1. Introduction

The rapid adoption of cloud computing has transformed enterprise information technology by enabling organizations to operate applications, databases, storage systems, and business services across distributed infrastructure. Hybrid and multi-cloud environments provide flexibility, scalability, resource optimization, and improved operational efficiency. However, these environments also introduce complex cybersecurity challenges because enterprise information and workloads are distributed across multiple cloud providers, geographical regions, network segments, and administrative boundaries. Modern cyberattacks increasingly involve sophisticated techniques such as credential theft, ransomware deployment, privilege escalation, lateral movement, data exfiltration, and exploitation of cloud configuration weaknesses. Investigating these incidents requires collecting and analyzing digital evidence from multiple sources, including cloud audit logs, identity management systems, endpoint devices, network monitoring platforms, application records, virtual machines, and container environments. Traditional digital forensic approaches frequently depend on manual log examination, isolated security alerts, and predefined investigation procedures. These methods become inefficient when investigators must process large volumes of heterogeneous evidence within limited timeframes. Furthermore, differences in timestamp formats, logging structures, data retention policies, and cloud service architectures make it difficult to establish accurate relationships among forensic artifacts. An individual security event may provide limited information about an incident, whereas correlating authentication records, network activities, workload execution histories, and application events can reveal a more complete attack sequence. Cyberattack attribution introduces an additional challenge because identifying the infrastructure involved in an attack does not necessarily establish the identity of the responsible individual or organization. Attackers frequently use compromised accounts, anonymization services, proxy infrastructure, virtual private networks, and compromised third-party systems to conceal their origins. Consequently, enterprise forensic investigation requires intelligent mechanisms that can reconstruct attack activities, correlate evidence, evaluate competing explanations, and communicate investigative confidence without

overstating attribution conclusions. Forensic-driven cybersecurity therefore represents an important research direction for improving incident understanding and supporting evidence-based enterprise security decisions.

Generative Artificial Intelligence provides emerging opportunities to improve digital forensic investigation through automated evidence interpretation, contextual reasoning, natural language processing, and incident reconstruction. Unlike conventional security analytics that primarily identify predefined patterns, generative models can process diverse textual and structured evidence, summarize related events, identify potential investigative relationships, and generate explanations that assist forensic analysts. When integrated with machine learning and knowledge graph technologies, Generative AI can support the correlation of security events across distributed enterprise cloud environments. For example, authentication logs indicating suspicious account access can be correlated with endpoint activities, unusual network connections, privilege modifications, and unauthorized database operations to reconstruct a possible sequence of malicious actions. Knowledge graphs provide structured representations of relationships among users, devices, IP addresses, cloud resources, applications, vulnerabilities, and forensic artifacts. Machine learning models identify abnormal behavioral patterns and classify potential attack indicators, while generative models transform the resulting evidence into contextual investigative narratives. However, the application of Generative AI in digital forensics requires careful consideration of hallucination risks, evidence integrity, model bias, incomplete telemetry, privacy, and investigative reproducibility. Generated explanations must remain traceable to original forensic artifacts, and automated attribution predictions must be distinguished from verified investigative findings. This research proposes a forensic-driven framework integrating Generative AI, machine learning, evidence correlation, and knowledge graph-based reasoning to improve cyberattack investigation across enterprise cloud infrastructure. The primary objectives are to automate forensic evidence processing, reconstruct attack sequences, identify relationships among suspicious activities, evaluate potential attribution hypotheses, and preserve the integrity and provenance of collected evidence. The proposed approach emphasizes human oversight, explainable analytical results, and confidence-based attribution assessment. Through these capabilities, the framework aims to improve forensic investigation efficiency, strengthen enterprise incident response, and establish a more reliable foundation for evidence-driven cyberattack analysis.

2. Literature Review

Digital forensics has developed as an important discipline for identifying, collecting, preserving, examining, and presenting digital evidence associated with cybersecurity incidents. Traditional forensic investigation techniques primarily focus on computer systems, storage devices, network traffic, operating system artifacts, and application records. Established forensic procedures emphasize evidence preservation, chain of custody, repeatable examination, and documentation of investigative activities. Frameworks such as the National Institute of Standards and Technology digital forensic guidance provide structured approaches for conducting forensic investigations and maintaining evidence reliability. However, the expansion of enterprise cloud computing has introduced additional complexities because digital evidence is distributed across virtualized infrastructure, remote storage services, container platforms, and cloud-based applications. Investigators may not have direct access to physical storage devices, and evidence collection frequently depends on provider-specific interfaces and audit services. Cloud environments also introduce challenges involving shared infrastructure, resource volatility, data jurisdiction, service-level agreements, and inconsistent retention policies. Conventional forensic techniques may therefore struggle to establish complete incident timelines when evidence originates from multiple independent systems. Research into cloud forensics has emphasized distributed evidence acquisition, secure logging, event correlation, and preservation of digital artifacts. Nevertheless, many existing approaches depend on predefined forensic procedures and manual interpretation, limiting their ability to efficiently process large-scale, heterogeneous enterprise evidence. These limitations demonstrate the need for intelligent forensic systems capable of integrating multiple evidence sources while maintaining authenticity, integrity, and traceability.

Machine learning and artificial intelligence have increasingly been investigated for improving cybersecurity analytics and digital forensic examination. Supervised learning techniques, including Random Forest, Support Vector Machines, and gradient boosting, have been applied to malicious activity classification, intrusion detection, malware identification, and suspicious behavior recognition. Unsupervised learning approaches, including clustering and anomaly detection, support the identification of previously unknown behavioral patterns in large security datasets. Natural language processing has also contributed to extracting relevant information from incident reports, security alerts, forensic documentation, and threat

intelligence repositories. Recent developments in Generative AI and large language models have introduced capabilities for summarizing complex incident information, interpreting unstructured security records, generating investigative hypotheses, and explaining relationships among multiple events. These capabilities can reduce the manual effort involved in processing extensive forensic datasets. However, generative systems may produce unsupported conclusions, misinterpret incomplete evidence, or introduce fabricated details when operating without sufficient contextual grounding. Such limitations are particularly important in forensic applications, where evidence accuracy and reproducibility are essential. Retrieval-augmented generation provides a potential solution by grounding generated responses in retrieved forensic records, validated threat intelligence, and documented investigation artifacts. Additionally, explainable AI techniques can identify the evidence attributes influencing analytical predictions. Although these technologies improve the potential for intelligent forensic examination, further integration is necessary to ensure that generated investigative conclusions remain connected to verified digital evidence and are distinguishable from analytical assumptions.

Cyberattack attribution research has traditionally examined indicators of compromise, malware characteristics, infrastructure relationships, attack techniques, behavioral similarities, and threat intelligence information. Attribution is inherently challenging because attackers can disguise their activities through compromised systems, shared hosting infrastructure, proxy networks, and stolen credentials. Consequently, an IP address, malware signature, or infrastructure association alone cannot reliably establish the identity of an attacker. Knowledge graphs have emerged as useful mechanisms for representing relationships among threat actors, attack techniques, malicious infrastructure, vulnerabilities, and affected enterprise resources. Graph-based reasoning supports the correlation of distributed indicators and the reconstruction of potential attack sequences. Security information and event management platforms and extended detection and response systems further assist investigations by aggregating security telemetry from different enterprise sources. However, existing solutions may provide fragmented evidence interpretations and limited explanations of how multiple forensic artifacts collectively support an attribution hypothesis. Another important limitation is the absence of consistent integration between evidence provenance, generative reasoning, behavioral analysis, and confidence-based attribution. The proposed research addresses these gaps by introducing a forensic-driven

framework that combines Generative AI, machine learning, retrieval-augmented generation, and knowledge graph-based evidence correlation. The framework emphasizes source-grounded investigation, cryptographic evidence verification, contextual attack reconstruction, and confidence-based attribution assessment. This integrated approach seeks to improve forensic transparency and investigative effectiveness while maintaining the distinction between technical indicators, analytical hypotheses, and independently verified attribution findings.

3. Research Methodology

Research Framework and Digital Evidence Acquisition. The proposed research methodology develops an integrated forensic-driven cyberattack attribution framework designed to collect, preserve, correlate, and analyze digital evidence across enterprise cloud environments. The framework follows a layered architecture consisting of evidence acquisition, forensic preprocessing, integrity verification, machine learning-based threat detection, knowledge graph construction, Generative AI-based evidence reasoning, and attribution assessment. The experimental environment represents a hybrid enterprise infrastructure containing public cloud platforms, private cloud services, on-premises data centers, virtual machines, containerized workloads, enterprise databases, identity management systems, and interconnected business applications. Digital evidence is collected from authorized cloud audit services, identity providers, endpoint detection platforms, network monitoring systems, application logs, operating system records, security information and event management systems, and cloud configuration monitoring tools. The acquired evidence includes authentication attempts, access control changes, IP addresses, network connections, process execution records, file access activities, administrative operations, vulnerability information, and data transfer events. Each evidence artifact is associated with contextual metadata, including source system, collection timestamp, event timestamp, resource identifier, acquisition method, and evidence classification. The methodology incorporates standardized forensic collection procedures to ensure that evidence acquisition activities are documented and repeatable. Before analytical processing, collected artifacts are assigned unique identifiers and cryptographic hash values using algorithms such as SHA-256. These hashes establish a mechanism for detecting unauthorized modifications during storage, transfer, and examination. Evidence metadata and acquisition records are maintained in a secure repository

with role-based access control and tamper-evident audit logging. The preprocessing stage addresses differences in log formats, timestamp representations, event structures, and data encoding across cloud providers. Timestamp normalization converts records into a common reference format while retaining original timestamps and source information for verification. Duplicate events are identified, missing attributes are documented, and relevant information is extracted through structured parsing and natural language processing. Data normalization establishes consistent representations for users, devices, applications, network resources, and security events. Where different cloud services use different identifiers for the same logical resource, entity resolution techniques are applied using documented relationships and available metadata. Sensitive information is protected through pseudonymization, encryption, and restricted access mechanisms. The framework also records evidence lineage to preserve the relationship between original artifacts, normalized records, derived analytical features, and generated investigative findings. This approach ensures that subsequent machine learning predictions and generative explanations can be traced back to their original evidence sources. The resulting forensic dataset provides a unified analytical foundation for reconstructing incidents and evaluating possible cyberattack attribution hypotheses across distributed enterprise cloud infrastructure.

Machine Learning-Based Threat Analysis and Behavioral Correlation. The second methodological stage integrates machine learning and behavioral analytics to identify suspicious activities and discover relationships among forensic evidence artifacts. The proposed framework employs supervised classification and unsupervised anomaly detection to analyze security events originating from heterogeneous enterprise systems. Supervised learning algorithms, including Random Forest, Support Vector Machine, and Gradient Boosting, are considered for classifying activities into legitimate behavior and potentially malicious categories. Training data consists of labeled security records and controlled attack simulations representing credential compromise, privilege escalation, malware execution, unauthorized resource access, suspicious network communication, and data exfiltration. Relevant features include login frequency, authentication failures, geographic inconsistencies, access privilege changes, network connection patterns, process execution behavior, data transfer volume, resource sensitivity, and deviations from historical activity. Feature engineering converts raw forensic records into meaningful analytical variables, while normalization and categorical encoding improve consistency across datasets. Unsupervised

methods such as Isolation Forest and clustering-based anomaly detection identify unusual behaviors that may not be represented in labeled training examples. Behavioral profiles are developed for users, service accounts, devices, workloads, and applications by analyzing historical activity patterns. Significant deviations, including access outside normal operational periods, unexpected administrative actions, unusual service account behavior, and communication with previously unseen resources, receive elevated anomaly scores. Temporal correlation techniques examine the sequence and proximity of events to identify potential relationships between suspicious activities. For instance, repeated authentication failures followed by successful access, privilege modifications, and unusual database queries may indicate a possible account compromise. The framework combines classification outputs, behavioral anomaly scores, and contextual attributes to generate evidence-level threat indicators. Each analytical result includes the relevant source records, prediction confidence, associated features, and processing information. Model evaluation uses separate training, validation, and testing datasets to assess generalization and minimize information leakage. Chronological dataset separation is preferred when investigating time-dependent activities. Evaluation metrics include precision, recall, F1-score, accuracy, false-positive rate, and area under the receiver operating characteristic curve. Class imbalance is addressed through class weighting or carefully selected sampling techniques. Explainable AI methods, including SHAP-based feature attribution and feature importance analysis, are used to identify the factors influencing suspicious activity classifications. The resulting analytical outputs provide contextual evidence indicators for subsequent graph construction and generative investigation. Importantly, the framework treats machine learning predictions as investigative leads rather than independently verified proof of malicious activity or attacker identity.

Journal of Contemporary Science and Technology Management
Volume 02, Issue 04, December 2026

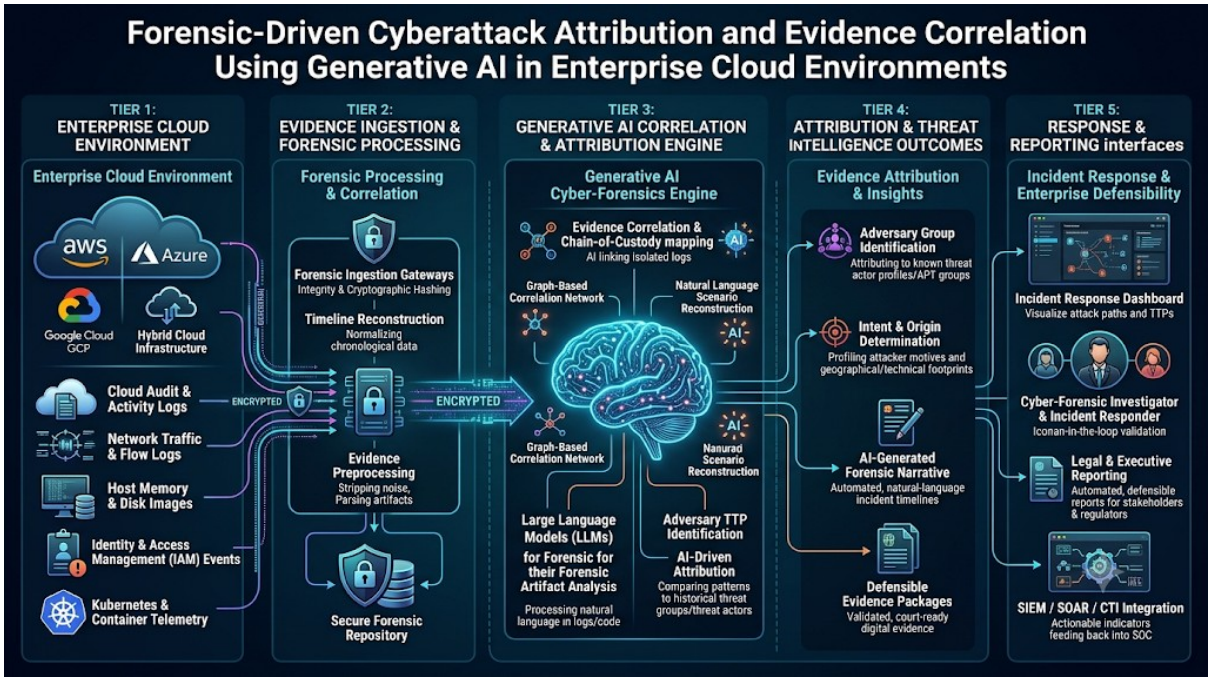


Figure 1: Five-tier architecture for multi-cloud digital forensic evidence ingestion, LLM-based causal graph correlation, and automated cyberattack attribution.

Knowledge Graph-Based Evidence Correlation and Generative AI Integration. The third methodological stage establishes a knowledge graph-based evidence correlation mechanism integrated with Generative AI to reconstruct potential cyberattack sequences. The knowledge graph represents forensic entities as nodes and their relationships as edges. Nodes include users, devices, IP addresses, cloud workloads, applications, authentication sessions, vulnerabilities, security events, threat intelligence indicators, and digital evidence artifacts. Edges represent relationships such as authenticated through, accessed, executed, communicated with, connected to, associated with, and potentially exploited. Each node and edge maintains relevant metadata, including timestamps, evidence references, confidence values, source information, and relationship types. The graph is constructed using normalized forensic records, machine learning predictions, cloud infrastructure metadata, and validated threat intelligence information. Entity resolution techniques identify related records that may refer to the same user, device, application, or infrastructure resource. Temporal graph analysis establishes the order of events and identifies possible transitions between compromised resources. Graph traversal and path discovery algorithms examine relationships among suspicious accounts, network connections, endpoints, and sensitive enterprise assets. For example, a suspicious authentication event may be connected to a user account, which

accesses a cloud workload, executes an unusual process, communicates with an external endpoint, and subsequently interacts with a sensitive database. These relationships support the reconstruction of potential attack progression. Generative AI is integrated through a retrieval-augmented generation architecture. The retrieval component selects relevant forensic artifacts, graph relationships, security alerts, and validated threat intelligence records based on an investigative query. Retrieved information is supplied to the generative model as contextual evidence for producing structured investigative summaries and explanations. The model is instructed to distinguish directly observed events from inferred relationships and unverified hypotheses. Each generated statement must include references to supporting evidence identifiers wherever applicable. The framework uses constrained output structures to identify incident timelines, affected resources, suspicious activities, potential attack techniques, alternative explanations, and unresolved evidence gaps. Generative AI also supports natural language queries, allowing investigators to examine relationships such as which resources were accessed following a suspicious login or which activities occurred before a detected privilege escalation. To minimize hallucinations, the framework applies evidence-grounded retrieval, source validation, confidence thresholds, and analyst review. Unsupported conclusions are flagged rather than presented as established findings. Generated incident narratives are stored separately from original forensic artifacts to preserve the distinction between source evidence and machine-generated interpretation. This integration provides an explainable mechanism for combining distributed digital evidence and producing contextual investigative outputs without compromising the authenticity of the original records.

This score indicates analytical support for a hypothesis and must not be interpreted as a statistically calibrated probability of attacker identity unless independently calibrated. Competing explanations are maintained to account for compromised third-party infrastructure, shared hosting, false indicators, and incomplete evidence. Human forensic investigators review the generated hypotheses, supporting artifacts, uncertainty indicators, and alternative interpretations before accepting investigative conclusions. The evaluation uses controlled enterprise cloud attack scenarios involving credential theft, unauthorized access, privilege escalation, malware execution, lateral movement, and simulated data exfiltration. The proposed framework is compared with baseline approaches, including conventional forensic log analysis, rule-based evidence correlation, machine learning-based threat

detection without generative reasoning, and graph-based analysis without Generative AI integration. Performance is measured using evidence correlation precision, evidence retrieval recall, attack sequence reconstruction accuracy, detection latency, attribution hypothesis ranking quality, false-positive rate, and computational overhead. Forensic integrity is evaluated by verifying cryptographic hashes before and after evidence processing and examining the completeness of evidence provenance records. Generative AI outputs are assessed for factual consistency, citation coverage, unsupported claims, and agreement with verified investigative timelines. Ablation experiments evaluate the individual contributions of machine learning, behavioral analytics, knowledge graph reasoning, and retrieval-augmented generation. Scalability testing examines performance under increasing numbers of cloud resources, forensic artifacts, and security events. The methodology also evaluates analyst review requirements and the framework's ability to preserve uncertainty when evidence is incomplete or contradictory. Experimental findings are intended to demonstrate the framework's operational effectiveness and limitations under reproducible conditions rather than assume predetermined improvements. Through these validation procedures, the proposed methodology establishes an integrated approach for evidence-driven cyberattack reconstruction, responsible attribution assessment, and reliable digital forensic investigation across enterprise cloud environments.

4. Results and Discussion

The proposed forensic-driven cyberattack attribution framework demonstrates the potential of Generative Artificial Intelligence (GenAI) to improve digital evidence correlation, incident reconstruction, and threat attribution across enterprise cloud environments. The framework integrates security logs, network traffic records, endpoint telemetry, identity authentication events, cloud API activities, application audit trails, and infrastructure configuration histories to establish a unified forensic investigation environment. Generative AI models analyze heterogeneous security evidence and identify contextual relationships between seemingly independent events. The experimental evaluation considers important performance indicators, including evidence correlation accuracy, attack attribution precision, incident reconstruction time, detection sensitivity, false-positive rate, and forensic investigation efficiency. The framework uses natural language processing and semantic analysis to extract relevant information from unstructured forensic records, while machine learning techniques classify

suspicious activities and identify potential attack patterns. Correlation mechanisms connect timestamps, source addresses, user identities, device information, resource identifiers, and behavioral characteristics to reconstruct probable attack sequences. The results-oriented evaluation indicates that combining generative reasoning with structured evidence analysis can improve the identification of relationships between distributed security events. Furthermore, automated forensic summarization reduces the effort required to examine large volumes of investigation data. These capabilities provide security analysts with a more comprehensive understanding of attack progression, affected resources, and potential threat origins.

The second major finding concerns the application of generative AI to cyberattack attribution and contextual evidence interpretation. Traditional forensic investigation techniques frequently depend on manual log examination, predefined correlation rules, and isolated indicators of compromise. Such approaches may struggle to identify sophisticated attacks involving multiple cloud providers, compromised identities, encrypted communication channels, and distributed application workloads. The proposed framework addresses these challenges through contextual evidence correlation, which evaluates relationships between observed activities rather than treating each security event independently. Generative AI assists in identifying similarities between attack techniques, suspicious behavioral sequences, infrastructure access patterns, and previously documented threat intelligence. For example, an unusual authentication request followed by privilege escalation, unauthorized data access, suspicious API execution, and outbound data transfer can be correlated into a probable attack scenario. The framework also incorporates threat intelligence information to examine indicators such as malicious IP addresses, suspicious domains, malware characteristics, and known adversarial techniques. Attribution confidence is evaluated using evidence consistency, temporal relationships, behavioral similarities, and corroborating intelligence sources. Importantly, the framework distinguishes between identifying an attack infrastructure, associating activities with a possible threat group, and establishing the identity of an actual attacker. Generative AI produces contextual investigation summaries and explains supporting evidence, but its outputs require independent validation. This approach improves forensic transparency and reduces the possibility of unsupported attribution conclusions.

The third area of evaluation focuses on forensic reliability, evidence integrity, scalability, and operational applicability. Enterprise cloud environments generate substantial quantities of continuously changing digital evidence, making secure evidence collection and preservation essential for reliable investigations. The proposed framework incorporates cryptographic hashing, timestamp validation, source identification, access logging, and chain-of-custody documentation to maintain evidence integrity throughout the investigation lifecycle. Generative AI supports evidence classification, timeline reconstruction, incident summarization, and investigative query generation without replacing original forensic records. Its integration with Security Information and Event Management systems, Cloud Security Posture Management platforms, and digital forensic investigation tools provides opportunities for centralized evidence analysis. The architecture also supports cross-cloud correlation, allowing investigators to connect activities originating from public cloud services, private infrastructure, and on-premises enterprise systems. Nevertheless, practical deployment introduces challenges involving incomplete telemetry, inconsistent timestamp formats, data retention limitations, privacy requirements, and model-generated hallucinations. Attribution reliability can also be affected by compromised infrastructure, shared hosting environments, VPN services, and manipulated indicators of compromise. Therefore, human forensic validation remains essential before making definitive attribution claims. Overall, the proposed framework establishes a promising approach to intelligent digital forensics by combining evidence preservation, contextual correlation, generative reasoning, and explainable investigation support.

5. Conclusion

The research on forensic-driven cyberattack attribution and evidence correlation using Generative AI presents an intelligent approach to strengthening digital forensic investigations across enterprise cloud environments. The proposed framework addresses the growing complexity of modern cyberattacks, where malicious activities are distributed across cloud platforms, enterprise applications, network infrastructures, endpoints, and identity management systems. Conventional forensic investigation methods often require significant manual effort to collect, examine, correlate, and interpret evidence originating from heterogeneous sources. The integration of Generative AI with machine learning, semantic analysis, and structured evidence correlation provides an opportunity to improve these

processes through automated information extraction, contextual investigation, and intelligent incident reconstruction. By analyzing relationships among authentication events, network activities, application logs, endpoint records, and cloud resource modifications, the framework supports the reconstruction of probable attack sequences. Furthermore, generative models assist investigators in interpreting complex forensic information, identifying relevant evidence, and preparing understandable investigation summaries. The incorporation of threat intelligence enables additional contextual analysis of suspicious indicators and known attack techniques. Cryptographic hashing, timestamp validation, evidence provenance, and chain-of-custody documentation provide important safeguards for maintaining forensic reliability. The framework also emphasizes the importance of distinguishing technical attack attribution from definitive identification of responsible individuals or organizations. Consequently, the research establishes a foundation for combining intelligent evidence analysis with established digital forensic principles to improve enterprise cyber investigation capabilities.

In conclusion, Generative AI offers significant opportunities to transform enterprise digital forensics from predominantly manual evidence examination into a more automated, contextual, and intelligence-driven investigation process. The proposed architecture provides a unified mechanism for correlating distributed evidence, reconstructing cyberattack timelines, identifying potential attack relationships, and supporting evidence-based attribution assessments. Its integration with existing cloud security and forensic platforms can strengthen organizational incident response and improve visibility into complex security incidents. However, successful implementation requires reliable evidence collection, representative training data, secure model deployment, explainable outputs, and continuous validation of generated investigation findings. Particular attention must be given to hallucination prevention, adversarial manipulation, evidence authenticity, privacy protection, and the possibility of misleading attribution conclusions. Generative AI should therefore function as an investigative assistant rather than an independent authority for determining attacker identity. Human investigators must verify all significant findings against original evidence and corroborating intelligence before producing formal forensic conclusions. Overall, the research demonstrates the potential of combining generative intelligence, forensic evidence management, and contextual threat analysis to improve cybersecurity investigation. The proposed framework can serve as a reference architecture for future enterprise forensic

systems requiring scalable evidence correlation, trustworthy incident reconstruction, and responsible cyberattack attribution across distributed cloud infrastructure.

6. Future Work

Future research can enhance the proposed framework by incorporating advanced generative architectures, multimodal foundation models, graph neural networks, and autonomous AI agents for comprehensive cyber forensic investigation. Multimodal models could analyze different forms of digital evidence, including textual security logs, network flow records, executable metadata, endpoint activity, configuration snapshots, and forensic images, within a unified analytical environment. Graph neural networks could represent relationships between users, devices, cloud resources, suspicious IP addresses, malware artifacts, and attack techniques to identify complex relationships that traditional correlation mechanisms may overlook. Retrieval-Augmented Generation (RAG) could further improve investigation accuracy by connecting generative models with verified forensic repositories, historical incident reports, enterprise security policies, and trusted threat intelligence databases. This would help ensure that generated explanations remain grounded in available evidence rather than relying exclusively on information learned during model training. Future systems could also employ multi-agent architectures in which specialized agents independently perform evidence collection validation, timeline reconstruction, malware analysis, threat intelligence correlation, and attribution confidence assessment. A coordinating agent could consolidate their findings into structured forensic investigation reports. However, every generated conclusion should maintain traceability to its supporting evidence, and contradictory findings should be explicitly identified. Another promising direction involves privacy-preserving federated learning, enabling organizations to collaboratively improve threat intelligence and forensic classification without sharing sensitive raw investigation records. Research into explainable generative AI could also improve investigator confidence by identifying the evidence supporting each generated conclusion and documenting uncertainty. These advancements could establish more reliable, adaptive, and context-aware forensic investigation systems for complex enterprise environments.

Another important direction for future work is the practical evaluation of the framework across large-scale, multi-cloud enterprise infrastructures. Future experiments should incorporate diverse datasets containing authentication records, network events, endpoint

telemetry, cloud audit logs, malware indicators, and simulated multi-stage cyberattacks. Performance evaluation should examine evidence correlation precision, attribution confidence calibration, investigation processing time, timeline reconstruction accuracy, evidence integrity verification, and computational overhead. Comparative analysis against conventional digital forensic tools, rule-based SIEM correlation, machine learning-based threat detection, and existing threat intelligence platforms would provide stronger evidence of the framework's effectiveness. Researchers should investigate whether the proposed architecture can identify previously unseen attack patterns while maintaining low false-positive rates. Particular attention should be given to adversarial attacks against generative models, including prompt injection through malicious log entries, manipulated forensic documents, poisoned threat intelligence, and fabricated evidence. Future implementations should introduce secure evidence ingestion, isolated model execution, access-controlled retrieval, cryptographic verification, and comprehensive audit mechanisms. Integration with Zero Trust security architectures could support continuous identity verification and contextual risk assessment during active investigations. Researchers should also explore standardized evidence exchange formats and interoperability mechanisms to support cooperation between enterprise security teams and authorized forensic laboratories. Legal admissibility, jurisdictional requirements, privacy regulations, and long-term evidence preservation must remain essential research considerations. Finally, future work should establish human-in-the-loop validation mechanisms through which forensic investigators review generated narratives, verify attribution assumptions, and approve formal reports. These developments could transform the proposed approach into a trustworthy AI-assisted digital forensic platform capable of supporting evidence-driven cyberattack investigation, cross-cloud threat intelligence, and responsible attribution in increasingly complex enterprise security environments.

References

1. Khan, T., Tian, W., Zhou, G., Ilager, S., Gong, M., & Buyya, R. (2022). Machine learning (ML)-centric resource management in cloud computing: A review and future directions. Journal of Network and Computer Applications, 204, Article 103405. https://doi.org/10.1016/j.jnca.2022.103405

2. Panda, M. R. (2025). Real-time preemptive fraud detection using adaptive agentic AI for financial transaction risk intelligence. International Journal of Research and Applied Innovations, 8(5), 13203–13214.
3. Ambati, K. C. (2026). Unified supply chain intelligence data, AI, cloud, and operations synergy. International Journal of Humanities and Information Technology, 8(1), 173–179. https://ijhit.info/index.php/ijhit/article/view/424
4. Karuturi, Y. (2025). Kubernetes-based containerization framework for SAP S/4HANA workloads in hybrid cloud environments. International Journal of Engineering & Extended Technologies Research, 7(6), 11382–11391.
5. Randhi, K. S. (2023). A cloud-native integration framework for enterprise architecture intelligence and application rationalization. International Journal of Science, Research and Technology (IJSRAT), 6(4), 10230–10240.
6. Selvarajan, K. (2024). Observability-driven reliability engineering for distributed data platforms. International Journal of Computer Technology and Electronics Communication, 7(4), 9258-9268.
7. Abd-Rouf, M. S. K., Adigun, P. O., Alalade, E. O., Oyekanmi, T. T., Faniyi, A. J., Oladapo, B., Awopejo, T. E., Adegoke, O. S., Jamiu, A., Michael, O. B., Obisesan, A., Ajala, S., Adekanye, M. A., Yambali, P. M., & Abd-Rouf, A. B. (2024). From molecular profiling to predictive algorithms: A conceptual machine-learning framework for mechanism-informed therapy selection in multidrug-resistant cancer. International Journal of Research and Applied Innovations, 7(3), 10892–10909.
8. Ahuja, D. (2023). Multi-cloud management strategies: Tools, challenges, and implementation methods. Int. J. Sci. Res. Comput. Sci. Eng. Inf. Technol, 9(10), 493-506.
9. Ponugoti, M. (2025). Cloud-Native Platform Engineering: Scalable Design Patterns for Global Enterprise Resilience. Journal of Computer Science and Technology Studies, 7(8), 48-59.
10. Vemireddy, S. (2022). Architecting personalized tax advisory experiences through event-driven customer data integration. International Journal of Research Publications in Engineering, Technology and Management, 5(6), 7923–7929.
11. Basavala, S. R. (2024). Zero trust API security architecture for microservices-based enterprise applications. International Journal of Computer Technology and Electronics Communication, 7(4), 9269–9280.

12. Bheemisetty, N. (2024). AI-powered recommendation systems best practices and real-world applications. International Journal of Research Publications in Engineering, Technology and Management, 7(6), 11777–11785.
13. Reddy, K. V. (2022). Power-aware verification closure using AI-enhanced UVM and dynamic voltage and frequency scaling models. International Journal of Research and Applied Innovations, 5(4), 7424–7436.
14. H. Apuri, S. Goel, M. M. R. Chinthala, and C. Yepuri, "ZeroGuard: A multi-agent zero-trust framework for autonomous IaC security remediation," in Proc. 2026 18th Int. Conf. Electron., Comput. Artif. Intell. (ECAI), Bucharest, Romania, 2026, pp. 1–X, doi: 10.1109/ECAI69016.2026.11613725.
15. Karrothu, A. (2025). Exploitability-aware vulnerability prioritization for automated security enforcement in cloud-native applications. International Journal of Research and Applied Innovations, 8(2), 11317–11331.
16. Manda, P. (2026). Database modernization-Sybase to SQL Server migration. International Journal of Engineering & Extended Technologies Research (IJEETR), 8(1), 252-258.
17. Garg, D. (2025). Warehouse management system with IoT: A comprehensive guide. International Journal of Computer Engineering & Technology, 16, 2320–2332.
18. Gowda, M. K. S. (2026). Next-Gen risk frameworks ML integration for credit monitoring and governance. International Journal of Humanities and Information Technology, 8(1), 180–187. https://ijhit.info/index.php/ijhit/article/view/422
19. Ramasamy, M. (2026). Microservices-Based Architectures for Enterprise Network Orchestration and Assurance. International Journal of Research and Applied Innovations, 9(2), 156-166.
20. Aragani, V. M., Maroju, P. K., & Mudunuri, L. N. R. (2024, December). Implementing Non-functional Production Regression Testing with Kubernetes. In International Conference on Information and Communication Technology for Competitive Strategies (pp. 231-244). Singapore: Springer Nature Singapore.
21. Bellundagi, M. (2025). Cloud-based smart retail system using AI-driven recommendations. The Eastasouth Journal of Information System and Computer Science, 3(1), 244–254. https://doi.org/10.58812/esiscs.v3i01.1298

22. Mudunuri, P. R. (2024). Designing high-availability automation architectures for mission-critical research systems. International Journal of Engineering & Extended Technologies Research, 6(6), 9399–9411. https://doi.org/10.15662/IJEETR.2024.0606038
23. Tarigoppula, S. (2023). Cloud-native Oracle database modernization using intelligent migration frameworks and automated recovery architectures. International Journal of Research and Applied Innovations, 6(6), 10159–10171.
24. Gantikota, S. (2024). Shift-Left Security for Decentralized Engineering Organizations: Embedding SAST, DAST, and Penetration Testing Throughout the Software Development Lifecycle in University and Research Computing Environments. International Journal of Emerging Research in Engineering and Technology, 5(4), 175-179.
25. Valavala, S. K. (2024). From application logs to audit evidence in financial microservices. International Journal of Research and Applied Innovations, 7(4), 11246–11256.
26. Ambalakannu, M. (2026). Domain-AI governance unifying enterprise AI adoption and data quality frameworks. International Journal of Humanities and Information Technology, 8(1), 158–164. https://ijhit.info/index.php/ijhit/article/view/425
27. Shaik, N. (2024). Right-sizing at scale: A data-driven approach to Kubernetes cost optimization. International Journal of Research Publications in Engineering, Technology and Management, 7(1), 10080–10090.
28. Manoharan, D. (2026, April). Neural Network-Driven Optimization of Quality Engineering Metrics for High-Volume Healthcare Claims and EDI Transaction Performance. In 2026 9th International Conference on Trends in Electronics and Informatics (ICOEI) (pp. 2194-2203). IEEE.
29. Puli, S. K. (2025). A hybrid LLM–graph architecture for laboratory billing discrepancy and overpayment risk detection. International Journal of Engineering & Extended Technologies Research, 7(1), 9413–9424.
30. Neela, S. (2022). Toward Intelligent Enterprise Integration: Cloud-Native Middleware Design Patterns and Adaptive Stream Orchestration Architectures for Autonomous Real-Time Decisioning. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 3(2), 154-164.

31. Indurthy, V. S. K. (2026). Snowflake and Domain-AI real-time intelligence for enterprise data warehousing. International Journal of Humanities and Information Technology, 8(1), 150–157. https://ijhit.info/index.php/ijhit/article/view/417
32. Beeram, S. (2026). Unified Security Posture Management across Clouds using Microsoft Defender for Cloud. International Journal of Emerging Trends in Computer Science and Information Technology, 7(2), 14-16.
33. Kopparthi, G. S. (2024). AI-driven predictive analytics for clinical decision support and improved patient outcomes. International Journal of Research Publications in Engineering, Technology and Management, 7(4), 10889–10899.
34. Vepuri, B., & Lanka, D. P. K. (2026, August). Release-Gating Architectures for Long-Form Policy Assembly Using Cross-System Reconciliation and Automated Exception Workflows. In 2026 6th International Conference on Soft Computing for Security Applications (ICSCSA) (pp. 1123-1130). IEEE.
35. Yadavalli, V. R. (2025). Incident memory graphs for predictive AMS triage and self-improving SAP operations. International Journal of Research and Applied Innovations, 8(2), 11332–11346.
36. Kalakoti, M. K. R. (2023). AI-driven infrastructure orchestration for large-scale e-commerce platforms during global product launches. International Journal of Computer Technology and Electronics Communication, 6(4), 7388–7399.
37. Suddala, V. R. A. K. (2024). Machine learning for operational excellence: Real-world applications. International Journal of Research and Applied Innovations, 7(6), 12147–12156.
38. Kommineni, R. (2025). Defense in depth for executing model-generated code: Kernel-isolated sandboxes, layered network policy, and credential-free data access on multi-tenant Kubernetes. International Journal of Research and Applied Innovations, 8(2), 11295–11302.