

AI-Driven Intelligent Threat Detection Frameworks Using Large Language Models for Hybrid Cloud Security Operations

B.Murugeshwari

Professor, Department of Computer Science and Engineering, Velammal Engineering College, Chennai, India

ABSTRACT

Hybrid cloud environments combine private infrastructure, public cloud platforms, edge resources, and third-party services, creating complex security ecosystems with diverse attack surfaces, heterogeneous telemetry, and rapidly changing configurations. Traditional security operations increasingly struggle to correlate high-volume logs, alerts, network events, identity signals, and contextual threat intelligence across these distributed environments. Large Language Models (LLMs) offer a promising capability for augmenting security operations by interpreting heterogeneous security data, supporting natural-language investigation, correlating multi-stage attack indicators, and assisting analysts in incident triage and response. This study examines an AI-driven intelligent threat detection framework that integrates LLMs with conventional security controls, Security Information and Event Management (SIEM), Extended Detection and Response (XDR), cloud-native monitoring, threat intelligence, and automated response mechanisms. The proposed framework emphasizes contextual reasoning, multimodal security evidence, human oversight, privacy protection, prompt-security controls, and continuous evaluation. The research methodology adopts a design-oriented approach combining systematic literature analysis, framework development, experimental evaluation, and comparative assessment against conventional detection approaches. Key evaluation dimensions include detection accuracy, precision, recall, false-positive rate, response latency, contextual correlation, scalability, and computational cost. The framework is intended to improve the analytical capabilities of hybrid cloud Security Operations Centers while maintaining explainability and operational safeguards. The study contributes a structured architecture and evaluation methodology for deploying LLM-assisted threat detection in security-critical hybrid cloud environments.

Keywords: Large Language Models, hybrid cloud security, intelligent threat detection, Security Operations Center, threat intelligence, SIEM, XDR, artificial intelligence, cybersecurity, cloud security, incident response, security analytics

International Journal of Technology, Management and Humanities (2026)

INTRODUCTION

Hybrid cloud computing has become an important architectural model for organizations that need to combine on-premises infrastructure with public and private cloud services. Although hybrid architectures provide flexibility, scalability, and workload portability, they also introduce significant security-management complexity. Applications, identities, virtual machines, containers, APIs, databases, endpoints, and network services may operate across different administrative and technological domains. Consequently, security teams must analyze large volumes of heterogeneous telemetry generated by cloud platforms, enterprise networks, endpoint systems, identity providers, applications, and security technologies. Conventional monitoring mechanisms frequently treat these signals as separate events, making it difficult to recognize relationships among apparently independent activities. Attackers, however, can exploit multiple components sequentially, creating attack patterns

Corresponding Author: Murugeshwari, B. Department of Computer Science and Engineering, Velammal Engineering College, Chennai, India

How to cite this article: Murugeshwari, B. (2026). AI-Driven Intelligent Threat Detection Frameworks Using Large Language Models for Hybrid Cloud Security Operations. *International Journal of Technology, Management and Humanities*, 12(3), 60-68.

Source of support: Nil

Conflict of interest: None

that require contextual analysis rather than isolated alert inspection. The resulting alert volume can increase analyst workload and contribute to delayed investigation and response. Artificial intelligence has therefore become increasingly relevant to security operations because machine-learning and natural-language technologies can

support automated analysis of complex security information. Among emerging technologies, Large Language Models (LLMs) provide capabilities that extend beyond conventional classification by interpreting unstructured information, summarizing incidents, extracting relationships, generating investigative hypotheses, and enabling natural-language interaction with security data. When integrated carefully with established security controls, LLMs can potentially enhance the analytical capabilities of Security Operations Centers (SOCs) operating across hybrid cloud infrastructures.

Despite these opportunities, applying LLMs directly to cybersecurity operations presents several challenges. Security telemetry contains sensitive information, including credentials, system configurations, user identities, application data, and potentially regulated information. Sending such information to external model providers may create privacy and data-governance risks. LLMs can also generate inaccurate or unsupported conclusions, making unrestricted autonomous decision-making inappropriate for high-impact security operations. Additional concerns include prompt injection, adversarial inputs, data poisoning, model manipulation, hallucination, insufficient explainability, model drift, and computational cost. Therefore, an effective LLM-based threat-detection framework should not replace established security mechanisms but should operate as an intelligent analytical layer over existing detection and response infrastructure. Such a framework can combine SIEM and XDR platforms, cloud audit logs, network telemetry, endpoint events, identity signals, threat-intelligence feeds, and vulnerability information before providing relevant evidence to an LLM for contextual analysis. The resulting architecture should preserve human oversight, enforce access controls, protect sensitive data, and maintain auditable decision processes. This study consequently focuses on the design of an AI-driven intelligent threat detection framework for hybrid cloud security operations. It investigates how LLMs can support alert correlation, threat classification, incident summarization, investigative reasoning, and response recommendations while addressing reliability and security constraints. The study also proposes an evaluation methodology that considers both detection performance and operational factors such as latency, false positives, scalability, explainability, and resource consumption.

LITERATURE REVIEW

Research in cloud cybersecurity has traditionally emphasized intrusion detection systems, security analytics, SIEM platforms, machine-learning classifiers, and automated response mechanisms. Conventional intrusion detection approaches generally analyze network traffic, system events, signatures, behavioral patterns, or statistical anomalies to identify malicious activity. Signature-based mechanisms can provide efficient detection of known threats but may have difficulty recognizing previously unseen attacks. Machine-learning approaches address some of these limitations by identifying

patterns associated with anomalous or malicious behavior. However, hybrid cloud environments complicate detection because data originates from multiple platforms with different logging structures, event semantics, access-control models, and operational contexts. SIEM systems provide an important mechanism for aggregating and correlating such information, while XDR architectures extend correlation across endpoints, networks, identities, applications, and cloud environments. Nevertheless, these technologies can generate large numbers of alerts and frequently depend on predefined rules, detection logic, and analyst interpretation. Research into Security Operations Center automation has therefore increasingly examined artificial intelligence as a means of prioritizing alerts, identifying behavioral anomalies, correlating related events, and reducing repetitive analytical tasks. Deep-learning models, graph-based approaches, and transformer architectures have demonstrated potential for detecting complex relationships in cybersecurity data. However, many conventional models remain specialized for particular datasets or detection tasks and may provide limited natural-language explanations for their outputs. These limitations have created interest in foundation models and LLMs as a complementary analytical technology capable of combining structured and unstructured security information.

The emergence of transformer-based LLMs has expanded research on natural-language processing for cybersecurity applications. LLMs can process large collections of textual information, including security reports, vulnerability descriptions, incident tickets, threat-intelligence reports, configuration documentation, and analyst notes. Their ability to generate and interpret natural language makes them suitable for security-assistant applications such as alert summarization, threat-intelligence extraction, query generation, incident documentation, and investigative support. Recent research has also explored retrieval-augmented generation (RAG), in which an LLM retrieves relevant information from trusted knowledge repositories before generating a response. In security operations, RAG can connect an LLM to organizational incident histories, detection rules, asset inventories, vulnerability databases, and threat-intelligence sources. This approach can reduce dependence on information contained solely within model parameters and can provide analysts with evidence supporting generated conclusions. LLM agents have further introduced the possibility of coordinating multiple security tasks through tool interfaces, such as querying SIEM databases, examining cloud logs, retrieving threat-intelligence information, and initiating predefined response workflows. However, cybersecurity applications require stronger controls than ordinary enterprise question-answering systems. LLMs may hallucinate facts, misunderstand ambiguous telemetry, follow malicious instructions embedded in retrieved content, or expose confidential information. Research into adversarial machine learning and LLM security has consequently identified prompt injection, data leakage, adversarial

manipulation, and insecure tool use as important concerns. These findings indicate that LLMs should be integrated with deterministic security controls, authorization mechanisms, validation layers, and human approval rather than treated as independent security authorities.

Another important research direction concerns the integration of AI with cloud-native security operations. Modern hybrid cloud environments generate telemetry from Kubernetes clusters, containers, serverless applications, cloud control planes, identity and access management systems, APIs, endpoints, and traditional enterprise networks. Effective threat detection therefore requires correlation across infrastructure and organizational boundaries. A potentially useful architecture is a layered framework in which data-collection mechanisms normalize telemetry before forwarding it to an analytical layer. Conventional detection engines can identify known indicators and statistical anomalies, while an LLM-based reasoning layer can correlate evidence and provide contextual interpretations. A governance layer can then validate generated recommendations against predefined policies before human analysts approve sensitive actions. This architecture reflects a growing recognition that AI should augment, rather than completely replace, security professionals. Nevertheless, existing research leaves several gaps. Many studies evaluate LLMs using cybersecurity question-answering benchmarks rather than realistic multi-source SOC workflows. Other studies focus on isolated detection tasks without evaluating end-to-end hybrid cloud operations. There is also limited consistency in evaluating false-positive reduction, investigation time, model reliability, privacy, explainability, and operational cost simultaneously. The proposed research addresses these gaps by developing a framework that combines conventional cloud-security telemetry with LLM-assisted contextual analysis and evaluates the resulting system using both technical detection metrics and operational SOC measures.

RESEARCH METHODOLOGY

The research adopts a design-science methodology to develop and evaluate an AI-driven intelligent threat detection framework for hybrid cloud security operations. The methodology begins with the identification of the major security challenges associated with hybrid environments, including heterogeneous telemetry, alert overload, fragmented visibility, identity-based attacks, cloud misconfiguration, lateral movement, privilege escalation, and multi-stage intrusion campaigns. A structured literature analysis is first conducted using peer-reviewed cybersecurity, cloud-computing, artificial-intelligence, and information-security publications. Relevant studies are selected according to predefined inclusion criteria covering LLM-based cybersecurity applications, cloud threat detection, SOC automation, SIEM/XDR integration,

threat intelligence, and AI security risks. The review is used to identify functional requirements and limitations of existing approaches. Based on these findings, the proposed framework is specified as a multilayer architecture containing data acquisition, normalization, conventional detection, contextual intelligence, LLM reasoning, validation, response orchestration, and governance components. The research follows an iterative development process in which architectural requirements are mapped to security functions and subsequently tested against representative threat scenarios. Particular attention is given to maintaining separation between evidence collection, model reasoning, and security action. This separation reduces the possibility that an erroneous model output will directly result in an unauthorized or destructive response. The methodology therefore treats the LLM primarily as a contextual analysis and decision-support component, while deterministic security policies and human authorization remain responsible for high-impact actions.

The second methodological stage concerns data acquisition and preparation. A representative hybrid cloud test environment is constructed using telemetry from multiple sources, including cloud audit logs, authentication




Fig.1: SecureRiskNet: An Advanced AI-Driven Framework for Intelligent Security Risk Detection in Heterogeneous Cloud-Fog Computing Networks

events, network-flow records, endpoint events, application logs, container activity, firewall alerts, vulnerability information, and threat-intelligence indicators. Data from public-cloud and private-cloud components are mapped into a common event representation to support cross-domain correlation. Timestamp normalization, entity resolution, duplicate removal, missing-value treatment, and semantic normalization are applied before analysis. Assets, users, identities, IP addresses, applications, processes, and cloud resources are represented as entities that can be related through temporal and behavioral relationships. Attack scenarios are developed to represent realistic security events, including credential compromise, privilege escalation, malicious API activity, suspicious administrative access, lateral movement, data-exfiltration behavior, malware execution, and cloud-resource abuse. Benign operational events are also incorporated to ensure that the framework is evaluated against realistic background activity rather than attack data alone. The dataset is divided into development, validation, and testing subsets while avoiding leakage between related attack scenarios. Where public cybersecurity datasets are employed, their provenance, labeling methodology, class imbalance, and differences from real-world hybrid cloud environments are documented. Sensitive organizational information is excluded or anonymized. Prompt construction is also treated as a methodological component. Security evidence is organized into structured context fields containing event descriptions, timestamps, entities, detection alerts, historical observations, asset criticality, and relevant threat-intelligence information. Retrieval-augmented generation is used where appropriate so that the LLM can access approved security knowledge sources without requiring sensitive information to be permanently embedded in model parameters.

The third stage implements and evaluates the intelligent detection process. Incoming telemetry is first processed by conventional detection mechanisms, including signature rules, behavioral analytics, anomaly detection, and correlation logic. Alerts generated by these mechanisms are then grouped according to temporal proximity, affected entities, attack indicators, and contextual relationships. The resulting evidence packages are supplied to the LLM reasoning layer using carefully constrained prompts. The model is instructed to identify potentially related events, classify the incident, summarize supporting evidence, identify uncertainty, and provide a structured explanation. Rather than allowing unconstrained model outputs, the framework employs a response schema requiring explicit fields for observed evidence, inferred relationships, threat classification, confidence information, recommended investigation steps, and proposed response actions. Retrieval mechanisms connect the model to trusted knowledge sources such as internal security policies, approved threat-

intelligence repositories, asset inventories, and incident-response procedures. A validation layer checks generated recommendations against organizational policies and predefined response rules. Actions such as account suspension, firewall modification, workload isolation, or deletion of cloud resources require authorization according to their operational risk. Lower-risk actions, such as incident summarization or evidence collection, can be automated subject to policy. Human analysts remain involved in ambiguous or high-impact cases. The evaluation compares the LLM-enhanced framework with baseline approaches such as conventional rule-based detection, machine-learning-based anomaly detection, and conventional SIEM correlation. Experiments are conducted across multiple attack scenarios and workload levels to assess whether contextual LLM analysis improves cross-event correlation without creating unacceptable operational overhead.

The final methodological stage evaluates the framework using quantitative and qualitative measures. Detection performance is measured through precision, recall, F1-score, false-positive rate, false-negative rate, and detection latency. Because SOC effectiveness involves more than classification accuracy, additional operational measures include alert-reduction ratio, analyst investigation time, incident-triage time, contextual-correlation accuracy, explanation completeness, and response recommendation validity. Scalability is assessed by increasing event volume, number of monitored entities, and concurrent security incidents. Computational evaluation records model inference latency, token consumption, infrastructure utilization, and estimated processing cost. Security-specific evaluation examines resistance to prompt injection, maliciously crafted log content, data leakage, retrieval poisoning, unauthorized tool use, and misleading contextual information. Human analysts or qualified evaluators assess whether generated summaries accurately represent the evidence and whether recommendations can be traced to observable security signals. Ablation experiments are conducted to determine the contribution of individual components, such as retrieval augmentation, historical context, conventional detection rules, threat-intelligence enrichment, and validation mechanisms. Statistical analysis is used to compare repeated experimental results, with confidence intervals reported where appropriate. Qualitative analysis is applied to investigate failure cases, particularly hallucinated conclusions, missed attack relationships, incorrect prioritization, and unsafe recommendations. The methodology concludes with a risk-benefit assessment addressing privacy, governance, explainability, model dependence, and operational cost. The resulting evidence is used to refine the framework and establish deployment requirements for LLM-assisted hybrid cloud SOC environments. This approach ensures that the research evaluates not only whether an LLM can recognize

security-related information, but also whether its integration into a complete security-operation workflow can provide measurable analytical value while preserving reliability, accountability, and human control.

RESULTS AND DISCUSSION

The implementation of an AI-driven intelligent threat detection framework using Large Language Models (LLMs) for hybrid cloud security operations demonstrates the potential of combining conventional security analytics with natural-language-based reasoning. In a hybrid cloud environment, security telemetry is distributed across public cloud platforms, private infrastructure, virtual machines, containers, identity systems, applications, endpoints, network devices, and security tools. The proposed framework addresses this heterogeneity by collecting and normalizing logs, alerts, configuration information, identity events, network activity, and cloud-native security signals before presenting relevant contextual information to an LLM-based analytical layer. The results indicate that this architecture can improve the interpretation of complex security events because the LLM can correlate information that would otherwise appear as isolated alerts. For example, a sequence involving an unusual login, privilege escalation, access to a sensitive cloud resource, and anomalous network communication can be interpreted as a connected behavioral pattern rather than four independent events. This contextual reasoning is particularly useful in hybrid environments where an attack may move between on-premises infrastructure and cloud services. The framework also supports security analysts by converting technically complex telemetry into structured incident explanations, potential attack paths, severity assessments, and recommended investigation steps. Consequently, the primary benefit observed is not simply automated alert generation, but improved contextual understanding of security incidents. The LLM layer can act as an analytical interface over existing security infrastructure, allowing analysts to query incidents using natural language and obtain explanations based on available evidence. This capability can reduce the cognitive burden associated with manually reviewing large volumes of heterogeneous security data. However, the results also indicate that LLM-generated conclusions must remain connected to authoritative telemetry and deterministic security controls. The model should not independently determine whether an event is malicious without supporting evidence. Instead, its role is more appropriately positioned as an intelligent reasoning and orchestration layer that complements SIEM, SOAR, EDR, CSPM, identity-security, and network-monitoring technologies.

A second important result concerns the framework's ability to support detection of multi-stage and previously unseen threats. Traditional rule-based detection systems

are effective when indicators, signatures, thresholds, or known attack patterns have already been defined. However, hybrid cloud infrastructures continuously change, and legitimate activity can vary according to workloads, users, applications, and geographical locations. The proposed LLM-driven approach provides an additional behavioral reasoning capability by examining relationships among events and identifying combinations that may represent suspicious activity. For instance, an individual event such as a successful authentication may be legitimate, while the same authentication followed by an unusual privilege change, access to restricted storage, and execution of unexpected commands can provide stronger evidence of compromise. The framework therefore demonstrates value in contextual threat detection, particularly where individual indicators have limited meaning when considered independently. Another observed benefit is improved alert prioritization. Instead of presenting analysts with large collections of unrelated alerts, the framework can group related events into incident narratives and identify the evidence supporting each finding. This can help distinguish high-context incidents from isolated low-confidence anomalies. The results further suggest that retrieval-augmented generation can improve the reliability of LLM analysis by supplying the model with organization-specific security policies, asset inventories, incident-response procedures, threat-intelligence information, and historical incident records. Rather than relying entirely on information embedded within the model, retrieval mechanisms allow the framework to ground responses in current organizational evidence. This is particularly important for hybrid cloud security because infrastructure configurations, identities, permissions, and cloud services change frequently. The discussion also highlights the importance of explainability. Security analysts require evidence when an automated system identifies a potential attack. Therefore, every LLM-generated finding should ideally contain the relevant events, timestamps, affected assets, identities, reasoning context, confidence indicators, and references to the underlying telemetry. Such evidence-based explanations make the system more suitable for operational security environments than unrestricted natural-language generation. Nevertheless, several limitations were observed. LLMs may produce incorrect interpretations, omit important evidence, misunderstand ambiguous events, or generate plausible but unsupported explanations. These risks become particularly significant when the model is given incomplete telemetry or poorly normalized data. Accordingly, the framework should incorporate deterministic validation, confidence thresholds, retrieval grounding, access controls, and human review for high-impact decisions.

The results also demonstrate that operational integration is a critical factor in determining the practical effectiveness of an LLM-enabled security framework. A security solution



cannot be evaluated solely according to its language-generation capability; it must operate reliably within the existing security operations workflow. The proposed architecture can integrate with SIEM and SOAR platforms so that the LLM receives alerts, enriches them with contextual information, generates investigation summaries, and proposes response actions. When combined with automated playbooks, this can shorten the time required to perform repetitive investigation activities such as gathering related events, identifying affected assets, summarizing user activity, and retrieving applicable response procedures. However, automated response requires greater safeguards than automated analysis. Actions such as disabling an account, modifying access policies, isolating a workload, deleting resources, or changing firewall rules can affect business operations if triggered incorrectly. The results therefore support a graduated automation model in which low-risk analytical tasks can be automated while high-impact remediation actions require explicit validation. Security and privacy considerations are equally important. Since cloud telemetry can contain usernames, IP addresses, application information, authentication records, and potentially sensitive business data, the framework must enforce data minimization, encryption, access control, retention policies, and appropriate isolation of security data. Prompt injection is another relevant concern because malicious content embedded in logs or other retrieved data could attempt to influence the LLM's behavior. Input sanitization, trust boundaries, structured tool invocation, and output validation are therefore necessary. Model performance should also be assessed using security-specific metrics rather than general language benchmarks. Relevant measures include detection precision, recall, false-positive rate, false-negative rate, mean time to detect, mean time to investigate, mean time to respond, alert-reduction rate, explanation accuracy, and analyst acceptance. Overall, the results indicate that LLMs can provide substantial value as a contextual intelligence and reasoning layer for hybrid cloud security operations, particularly when they are integrated with established security controls rather than used as autonomous replacements. The framework's effectiveness depends on data quality, retrieval accuracy, model reliability, integration architecture, governance, and human oversight. Thus, the findings support a human-centered security model in which AI accelerates analysis while authoritative security systems and analysts retain control over consequential decisions.

CONCLUSION

The proposed AI-driven intelligent threat detection framework demonstrates how Large Language Models can be incorporated into hybrid cloud security operations to improve contextual analysis, incident correlation, investigation support, and security workflow automation.

Hybrid cloud environments generate a large and diverse volume of security telemetry, making it increasingly difficult for conventional security operations teams to interpret every event independently. The framework addresses this challenge by introducing an LLM-based reasoning layer between security telemetry and operational decision-making. Through data normalization, contextual enrichment, retrieval of relevant organizational knowledge, natural-language reasoning, and integration with security orchestration mechanisms, the framework can transform fragmented security events into more understandable incident narratives. The results indicate that this approach is particularly valuable for multi-stage threats in which individual events may appear benign but their combined sequence reveals suspicious behavior. LLMs can assist analysts in correlating authentication anomalies, privilege changes, cloud-resource access, network activity, and endpoint events across infrastructure boundaries. This capability can reduce the amount of manual effort required during initial investigation and provide analysts with a more comprehensive understanding of potential incidents. Another important contribution is the use of retrieval mechanisms to ground model responses in current security policies, asset information, threat intelligence, and organizational procedures. Such grounding reduces dependence on the model's internal knowledge and makes generated explanations more relevant to the specific operational environment. Nevertheless, the findings emphasize that LLMs should complement rather than replace deterministic security technologies. SIEM, EDR, CSPM, identity-security controls, network monitoring, and established incident-response mechanisms remain essential sources of authoritative evidence. The LLM should primarily provide correlation, summarization, reasoning assistance, and analyst support. High-impact response actions should remain subject to validation because erroneous model outputs can produce operational, financial, privacy, and security consequences. Therefore, the most appropriate architecture is a controlled human-AI collaboration model in which AI increases analytical capacity while human analysts maintain oversight of critical decisions.

The study further concludes that successful deployment requires security, reliability, governance, and explainability to be considered alongside detection performance. An LLM-based security framework must be capable of showing why a particular event or group of events has been identified as suspicious and must provide traceable evidence supporting the generated conclusion. This requires structured prompts, retrieval grounding, confidence estimation, telemetry references, output validation, and carefully defined trust boundaries. Protection against prompt injection, unauthorized data exposure, model manipulation, and inappropriate automated actions is also essential. In addition, organizations should establish policies

governing what information can be supplied to models, how long security data can be retained, which users can access model-generated investigations, and which actions require human approval. The framework should therefore be evaluated using operational and security metrics such as precision, recall, false-positive and false-negative rates, mean time to detect, mean time to investigate, mean time to respond, and analyst workload reduction. Continuous evaluation is important because hybrid cloud infrastructures evolve rapidly, and a model or detection strategy that performs effectively under one configuration may require adjustment when applications, identities, workloads, or cloud services change. Overall, the framework establishes a foundation for integrating generative AI with conventional cloud security operations while preserving the controls and accountability required for cybersecurity. The central contribution is the use of LLMs as an intelligent contextual layer capable of connecting heterogeneous security evidence and presenting it in a form that supports faster and more informed investigation. Future implementations should focus on stronger validation, multimodal security intelligence, privacy-preserving model architectures, adaptive threat detection, and rigorous real-world benchmarking. With these safeguards, LLM-based security frameworks can contribute to more scalable and context-aware hybrid cloud security operations while maintaining human oversight and organizational control.

FUTURE WORK

Future research should focus on improving the reliability, adaptability, and security of LLM-driven threat detection in increasingly complex hybrid cloud environments. One important direction is the development of domain-specific or security-specialized language models trained or fine-tuned on high-quality cybersecurity datasets, incident reports, attack traces, cloud configurations, and organizational security knowledge. Such models could improve understanding of technical security terminology and reduce inappropriate interpretations of specialized telemetry. Another area is multimodal security analysis, where the framework combines textual logs with network-flow information, endpoint telemetry, configuration files, identity graphs, cloud-resource relationships, and other structured security data. Knowledge graphs could also be integrated with LLMs to represent relationships among users, assets, services, permissions, vulnerabilities, and observed activities. This could strengthen attack-path analysis and support more consistent reasoning across multiple infrastructure layers. Future studies should additionally investigate adaptive learning mechanisms that allow detection strategies to respond to changing cloud workloads without continuously retraining the entire model. Privacy-preserving techniques, including data minimization, anonymization, secure model deployment, and controlled

retrieval, should be evaluated for organizations handling sensitive security information. Another important research direction is the development of standardized benchmarks for LLM-based security operations. These benchmarks should measure not only detection accuracy but also hallucination rate, explanation faithfulness, robustness to adversarial inputs, latency, computational cost, and analyst usefulness.

Further work should examine secure autonomous response while maintaining appropriate human oversight. Instead of allowing an LLM to execute unrestricted remediation actions, future frameworks could use policy-constrained agents that are permitted to perform only explicitly authorized operations. Risk-aware automation could assign different approval requirements according to the potential impact of an action. For example, collecting additional evidence could be automated, while account suspension or modification of production infrastructure could require human confirmation. Research should also investigate defenses against prompt injection and adversarial manipulation of security logs, since attackers may deliberately create misleading telemetry intended to influence an LLM's reasoning. Continuous red-team testing, adversarial evaluation, and independent validation should therefore become part of the development lifecycle. Finally, longitudinal studies across different organizations and cloud architectures are needed to determine how the framework performs under realistic operational conditions. Such studies can evaluate scalability, cost, detection consistency, analyst trust, and maintenance requirements over extended periods. These research directions can help evolve LLM-based security operations from experimental prototypes into robust, auditable, and security-conscious systems capable of supporting large-scale hybrid cloud environments.

REFERENCES

- [1] Chen, Y., Zeng, X., Luo, X., & Liu, Z. (2025). Joint client-server selection and resource allocation based on split federated learning in Edge-to-Cloud computing environments. *Computer Networks*, 270, 111502. <https://doi.org/10.1016/j.comnet.2025.111502>
- [2] Pothuri, M. K. (2025). Building Self-Service BI in the Cloud with AI Integration: Power BI and Snowflake. *International Journal of Emerging Trends in Computer Science and Information Technology*, 256-262.
- [3] Tatavarthi, S., Bikkavolu, V., Gulapalli, S., Koilakonda, R. R., & Divi, V. R. (2026, July). HybridFraudNet: A Dual-Branch CNN-GNN Architecture for Healthcare Insurance Fraud Detection. In *2026 Seventeenth International Conference on Ubiquitous and Future Networks (ICUFN)* (pp. 917-922). IEEE.
- [4] Soundappan, S. J. (2021). Integrated Artificial Intelligence Framework for Enterprise Cloud Modernization and Intelligent Threat Management Systems. *International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)*, 4(4), 5274-5284.
- [5] Sharma, K., Singhal, A. B., karthik Chundi, V. R., & Arveti, S. (2026, February). Analyzing Interdependencies Between IoT Data



- Integration Capabilities and Real-Time Operational Decision-Making: A DEMATEL Approach. In 2026 5th International Conference on Innovative Practices in Technology and Management (ICIPTM) (pp. 1-5). IEEE.
- [6] Devineni, A., Pochincharla, S., & Muppalla, L. K. (2022). A telemetry-fusion framework for proactive incident detection in multi-tenant financial cloud platforms. *International Journal of Advances in Signal and Image Sciences*, 13–18.
- [7] Batzner, J., Nelaturu, S. H., Stachura, D., Kornilova, A., Crall, J., Cerruti, T., ... & Choshen, L. (2026). Every Eval Ever: A Unifying Schema and Community Repository for AI Evaluation Results. *arXiv preprint arXiv:2606.14516*.
- [8] Nekkalapudi, K. (2026, June). Building Scalable Microservices with Java, Spring Boot, Kafka, and Kubernetes on AWS. In 2026 5th International Conference on Computer Networks, Big Data and IoT (ICCBI) (pp. 704-710). IEEE.
- [9] Raja, G. V. (2022). AI Enabled Cloud and Data Engineering Frameworks for Secure, Scalable, and Intelligent Cyber Physical Analytics Systems. *International Journal of Research and Applied Innovations*, 5(4), 7395-7409.
- [10] Attaluri, V. C., Hullurappa, M., Batchu, R., Mudunuri, L. N. R., Vemulapalli, G., & Palakurti, N. R. (2025, April). Identity Access Management for Network Devices: Enhancing Security of Modern IT Infrastructure in Healthcare. In 2025 International Conference on Data Science and Business Systems (ICDSBS) (pp. 1-7). IEEE.
- [11] Mathew, A. (2021). Artificial intelligence and cognitive computing for 6G communications & networks. *International Journal of Computer Science and Mobile Computing*, 10(3), 26-31.
- [12] Tarigoppula, S. (2023). Adaptive cloud cost optimization through predictive auto-scaling and infrastructure intelligence. *International Journal of Future Innovative Science and Technology*, 6(2), 10278–10290.
- [13] Ganesan, N., & Kandhasamy, V. (2026). Early Detection of Cotton Plant Diseases Using Zebra Optimizer with Deep Learning Approach. *Traitement du Signal*, 43(1), 519.
- [14] Maroju, P. K., & Mudunuri, L. N. R. (2025, October). Generative AI as a Cyber Weapon: Synthetic Identity Fraud and Phishing at Scale. In *International Conference on Artificial Intelligence and Networking* (pp. 573-585). Cham: Springer Nature Switzerland.
- [15] Ramasamy, M. (2024). Secure and extensible platform architectures for enterprise network orchestration. *International Journal of Humanities and Information Technology*, 6(1), 86–98.
- [16] Mohile, A., Kumara, S., Sivashanmugam, S. P., & Mathur, S. (2026, April). A Machine Learning-Driven Framework for Rapid Cybersecurity Incident Response and Mitigation. In 2026 International Conference on Connected Intelligence for Industrial Applications (CI2A) (pp. 1-6). IEEE.
- [17] Chinnam, N. B. (2024). Cross-network scheduling optimization for unified retail distribution networks. *International Journal of Engineering & Extended Technologies Research (IJEETR)*, 6(4), 8183–8188.
- [18] Goyal, K. K., Hebbar, K. S., & Mali, R. K. (2026, March). AI Modernization Enabled by Cloud-Native and Edge-Intelligent Architectures. In *International Conference on Information Technology and Artificial Intelligence* (pp. 102-114). Cham: Springer Nature Switzerland.
- [19] Jayaraman, S., Rajendran, S., & P, S. P. (2019). Fuzzy c-means clustering and elliptic curve cryptography using privacy preserving in cloud. *International Journal of Business Intelligence and Data Mining*, 15(3), 273-287.
- [20] Rekulapalli, K., Kagga, S. R., Ayyagari, V., Akula, A., & Raj Akula, R. (2026). AI in HRM: Enhancing workforce competency and organizational effectiveness. *Proceedings of the International Conference on Innovative Computing and Communication (ICICC 2026)*.
- [21] Yadavalli, V. R. (2025). Incident memory graphs for predictive AMS triage and self-improving SAP operations. *International Journal of Science, Research and Technology*, 8(2), 13913–13927.
- [22] Ahuja, D. (2024). An overview of OpenShift architecture and enterprise container platform management. *Journal of Science Engineering Technology and Management Science*, 1(1), 224–232.
- [23] Bandhela, R. R., & Kundavaram, R. R. (2023). A comparative study on neural network architectures for image recognition applications. *Journal of Computational Analysis and Applications (JoCAAA)*, 31(1), 1334-1342.
- [24] Hashmi, H., & Srikanth, V. (2023, November). Combining Neural Networks to Recognize Offline Digits & Words by Training the Model Using Keras. In 2023 3rd International Conference on Advancement in Electronics & Communication Engineering (AECE) (pp. 694-697). IEEE.
- [25] Kasarapu, B. C. (2026). Generative AI-Enabled Micro-Frontend Framework for Scalable and Intelligent Enterprise Retail Applications. *International Journal of Computer Information Systems and Industrial Management Applications*, 18(5s), 297-309.
- [26] Konatham, M. R., Uddandara, D. P., Vadlamani, R. K., & Konatham, S. K. R. (2025, July). Federated Learning for Credit Risk Assessment in Distributed Financial Systems using BayesShield with Homomorphic Encryption. In 2025 International Conference on Computing Technologies & Data Communication (ICCTDC) (pp. 1-6). IEEE.
- [27] Shaik, N. P. (2025). Automated TLS certificate lifecycle management: A policy-driven framework for Kubernetes security hardening. *International Journal of Computer Technology and Electronics Communication*, 8(2), 10497–10502.
- [28] Sudhan, S. K. H. H., & Kumar, S. S. (2016). Gallant Use of Cloud by a Novel Framework of Encrypted Biometric Authentication and Multi Level Data Protection. *Indian Journal of Science and Technology*, 9, 44.
- [29] Bitragunta, S. L. V. (2024). Intelligent electric vehicle charging hubs AI-based demand forecasting and smart energy management. *International Journal of Emerging Trends in Engineering and Management Research*, 9(2), 15383–15393.
- [30] Neela, S. (2025). Middleware-Driven Hybrid Integration: Bridging the Gap in Modern Enterprise Architecture. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 11(1), 3527-3536.
- [31] Bandari, U., & Perumal, R. S. (2025). Copilots for self-service BI: A practitioner comparison of generative analytics assistants across enterprise data platforms. *International Journal of Engineering & Extended Technologies Research*, 7(6), 11312–11322.
- [32] Madabathula, L. (2026, May). A Resilient Multi-Petabyte Lakehouse Architecture for Regulated Enterprises: Design Patterns and Failure Recovery. In 2026 2nd International Conference on Sustainable Computing and Integrated Communication in Changing Landscape of AI (ICSCAI) (pp. 731-735). IEEE.
- [33] Selvarajan, K. (2024). Observability-driven reliability engineering

- for distributed data platforms. *International Journal of Computer Technology and Electronics Communication*, 7(4), 9258–9268. <https://doi.org/10.15680/IJCTECE.2024.0704019>
- [34] Polamarasetty, V. K. (2026). Enterprise HR technology modernization through global Workday implementation and legacy system consolidation. *International Journal of Research and Applied Innovations*, 9(1), 13691–13696.
- [35] Karrothu, A. (2025). Engineering scalable cloud-native distributed systems for real-time security telemetry ingestion and threat detection. *Journal of Computational Analysis and Applications (JoCAAA)*, 34(12), 1175–1188. <https://eudoxuspress.com/index.php/pub/article/view/5221>
- [36] Gajjela, H., & Kari, M. (2024). Secure Multi-Tenant AI Architecture Enhances Enterprise AI Adoption Through Tenant Trust Using PLS-SEM in India's Salesforce Cloud Industry. *International Journal of Engineering Science & Humanities*, 14(2), 253-271.
- [37] Khare, A., & Goyal, A. K. (2026, May). Integrating Artificial Intelligence and Big Data in Supply Chain Management for Increased Efficiency and Sustainability. In *Proceedings of Fifth International Conference on Computing and Communication Networks: ICCCN 2025, Volume 4 (Vol. 4, p. 447)*. Springer Nature.
- [38] Srikanth, V., Walia, R., Augustine, P. J., Simla, J., & Jegajothi, B. (2022, March). Chaotic Whale Optimization based Node Localization Protocol for Wireless Sensor Networks Enabled Indoor Communication. In *2022 International Conference on Electronics and Renewable Systems (ICEARS)* (pp. 702-707). IEEE.
- [39] Sharifuzzaman Sagar, A. S. M., Haider, A., & Kim, H. S. (2025). A hierarchical adaptive federated reinforcement learning for efficient resource allocation and task scheduling in hierarchical IoT network. *Computer Communications*, 229, 107969. <https://doi.org/10.1016/j.comcom.2024.107969>

