International Journal of Applied Mathematics, Sciences, and Technology for National Defense
Volume 3, Issue 3, 143-152
e_ISSN: 2985-9352, p_ISSN: 2986-0776
https://www.journal.foundae.com/index.php/JAS-ND
DOI : https://doi.org/10.58524/app.sci.def.v3i3.670

Next-Gen SOC: Leveraging generative AI for scalable threat detection and AI-powered alert classification

Sudheer Kotilingala*
IBM Corporation
USA

Article Info

Article history:

Received: May 21, 2025
Revised: August 12, 2025
Accepted: September 2, 2025
Published: December 30, 2025

Keywords:

Artificial Intelligence
Large Language Models
Security Threats
SIEM
SOC

Abstract

The volume of alerts produced by the SIEM system causes SOC analysts to experience alert fatigue, with actual security incidents obscured by more than fifty percent of notifications being considered false positives. This inefficiency causes delays in response times and puts organisations at risk due to insufficient resource allocation. We have, therefore, introduced a new framework in this paper, which incorporates LLMs into SOC initiatives. Overall, with the help of contextual understanding elements of LLMs, our framework concludes with 95,5% accuracy to classify the alerts as false positives or actual threats. The study's results, therefore, validate less alert fatigue, improved systems functioning, and shorter time to critical security events using the proposed methodology. As a result, this paper outlines the proposed system's description, development, and evaluation to determine its potential for future SOC operations.

To cite this article: Kotilingala, S. (2025). Next-Gen SOC: Leveraging generative AI for scalable threat detection and AI-powered alert classification. *International Journal of Applied Mathematics, Sciences, and Technology for National Defense,* 3(3), 143-152

INTRODUCTION

The abundance of data currently produced in today's dynamic threat environment means SOCs are overwhelmed by the number of alerts that originate from SIEM systems. The problem is identifying the noisy threats from the actual ones, which wastes time. To this end, new AI-based technologies that could improve SOC work and accelerate threat identification are being considered.

Current state Security Information and Event Management (SIEM) systems produce millions of daily alerts, including informational, warning, and alert levels. It's important to realise that while these systems are designed to err on the conservative side, this leads to too many false alarms. SOC analysts working on these solutions often waste up to fifty per cent of their time responding to false-positive and low-priority alerts. Alarm fatigue delays promoting existing threats and simultaneously diminishes SOC team efficiency and psychological well-being. The modern advanced threats and developing cybercriminal techniques create devastating outcomes when there are procrastinations in threat recognition and response (Ali et al., 2025).

The primary purpose of this research work is to fill the gap within SOC's existing traditional work processes and bring in advanced AI technology to enhance its operations. Due to their high NLP capabilities, Large Language Models (LLMs) present an excellent chance for changing the alert classification and SOC picture (Reddy & Reddy, 2024). The objectives of this research are to develop a scalable and context-aware alert classification framework, minimise false-positive rates while maintaining high accuracy in threat detection, reduce the workload of SOC analysts, allowing them

***Corresponding Author:**
Sudheer Kotilingala, IBM Corporation, USA, Email: reachsudheer.kotilingala@gmail.com

to focus on high-priority alerts, and integrate LLMs seamlessly into existing SIEM and SOAR (Security Orchestration, Automation, and Response) systems.

Related Work

Scholars have attempted to capture the development of Security Information and Event Management (SIEM) over the past decade to address the increasing complexity of threats. This section also discusses current and recent methods and techniques in the field, their capabilities, weaknesses, and scope for enhancement.

Existing SIEM approaches

Embedded within traditional SIEM systems exist predefined rule-based systems that produce alerts. The conventional detection methods succeed with known threats yet successfully detect new attack forms poorly, leading security operators to receive numerous incorrect alerts. The primary detection technique of SIEM systems includes heuristic-based analysis of specific IP addresses and port activities, which produces too many false alerts that burden SOC teams (Reddy & Reddy, 2024). The adoption of SIEM models enhanced through machine learning techniques has become a solution to improve detection performance and categorisation by implementing decision trees, deep neural networks, and random forests (Zeinali, 2016). ML-based SIEM systems encounter weaknesses: inadequate training data quality, labeling errors, and inability to monitor comprehensive attack scenarios, resulting in incorrect benign-vs-malicious identifications (Ferrag et al., 2020; Hassanin & Moustafa, 2024). The current approach involves a combination of AI models which unite rule-based logic frameworks with anomaly detection systems generated through ML, creating more accurate SIEM platforms (González-Granadillo et al., 2021).

Advances in AI and NLP models for SIEM

Modern AI-driven cybersecurity tools leverage large language models (LLMs) and NLP algorithms to analyse unstructured log data, security event records, and analyst notes (Moustafa & Slay, 2016; Shaukat et al., 2020). Generative AI models, such as GPT-based systems, have efficiently classified security incidents, summarised alerts, and automated incident response workflows (Brown et al., 2020). One promising AI technique in SIEM is few-shot learning, where models classify threats based on limited labeled samples (Radford et al., 2018). This is particularly useful in SOC environments, where labeled attack datasets are scarce (Yadav & Rao, 2015). However, NLP-powered SOC solutions also raise concerns regarding data privacy, real-time threat evaluation, and adaptability to evolving cyber threats (Goodfellow et al., 2020). Research indicates ways to protect data privacy through AI techniques despite using AI to fight security threats (Kshetri, 2025).

XAI and hybrid AI models for threat detection

ML-based SIEM systems encounter a significant hurdle when users cannot understand how the models reach their outcomes. Traditional deep learning models operate as black boxes, making it difficult for SOC analysts to interpret alert classifications and threat intelligence insights (Lundberg & Lee, 2017). Explainable AI (XAI) techniques, such as decision trees and SHAP (SHapley Additive Explanations), are increasingly incorporated into SIEM tools to provide transparent, human-readable threat insights. New research shows that combining rule-based heuristics with ML-based anomaly detection performs well in detecting threats and provides better explanations (Alwarafy et al., 2020). The correlation of various threat indicators between hybrid AI solutions composed of graph-based anomaly detection and ensemble learning models substantially enhances SOC performance.

Blockchain-SIEM integration for enhanced security

Researchers study blockchain technology as an answer to enhance SIEM performance through secure log storage alongside the decentralisation of security data sharing and fraud defense capabilities (Ferrag et al., 2024; Salah et al., 2019). The BlockSIEM implementation of blockchain technology provides security log storage with unalterable features to protect integrity throughout cyberattacks (Sun, 2020). This is particularly useful in smart city infrastructures, cloud-based SOCs, and IoT security frameworks. Researchers suggest integrating blockchain with AI-driven SIEM can further enhance threat intelligence, automate trust validation, and improve threat response mechanisms (Hyun et al., 2019; Liang et al., 2017). Many issues, like system growth limits and high computer usage, still block the widespread use of this technology (Pathak et al., 2024).

METHOD

Architecture Overview

The overall system analysis shows how separate parts and sub-components combine to perform security actions and identify risks at different times.

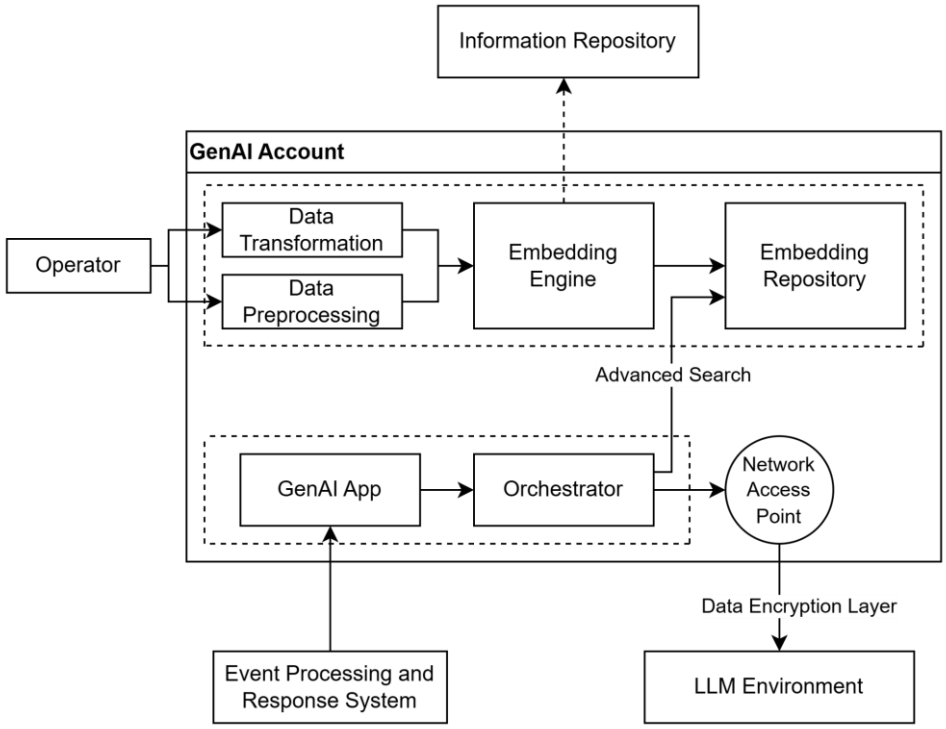

Figure 1. System Architecture Overview

Figure 1 shows how this solution operates by handling information events, transforming data, collecting data through SIEM, and creating vector embeddings before saving them to storage. They schedule the processes, do that themselves, and relay the alerts to the LLM, which is further hosted for contextual interpretation and categorization. Last, these are the validated alerts sent to the SOAR platform to comprehend the automated response workflows and deal with the incidents efficiently. The advanced technologies and secure data flows are illustrated with the help of diagrams as integral flows in architecture.

Ingesting SIEM data for real-time monitoring

The architecture starts by consuming logs and alerts from raw and unstructured source data, which is then collected and aggregated from other systems utilizing the SIEM system. The Rules Engine or Analytics Engine originates these events: logs of firewalls, IDSs, and network devices. The ordered events become the primary data source for processing during the subsequent analysis stages. This step provides continuous real-time security monitoring to watch everything happening in our environment of interest.

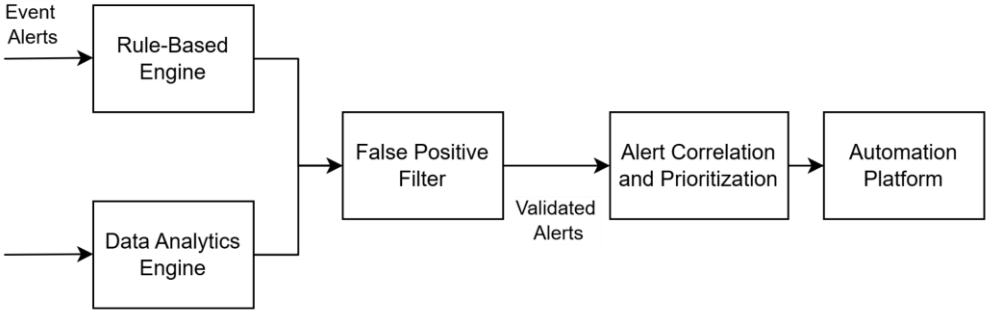

Figure 2. SIEM alerting workflow

Figure 2 mainly shows the alerting aspect of the SIEM system. It describes how the real-time alerts are created through the rules-based or advanced analytics alerting engines, then pass through the first filter via a fine-tuned models (FM) false positive detector module. These accurate positive alerts are further ranked and enriched before being forwarded to the SOAR platform, which is automated to deal with the incident. Some of the components of the proposed workflow are the non-creation of false positives, prioritization of alerts and optimization of the work of SOC analysts.

Pre-processing for alert enrichment

After events have been ingested, they go through a stringent pre-processing stage. The raw data collected is cleaned up and put into the appropriate structures to fit in specific data sources. Many duplicated events are redundant or repetitive to reduce the load for further processing, and other ambient noise is prevented from becoming part of the analysis. Further, the pre-processing step provides contextual and time-based metadata information, including timestamps, geographic info, device details and IP mapping.

Constructing semantic embeddings

The following action involves building semantic embeddings for all processed alerts. Higher-level methodologies such as transformers make these embeddings to capture the context and connection within the data. These semantic representations are stored in an arbitrary vector database, including Pinecone or Weaviate, for efficient search during semantic queries. It positively affects vector identification of the various similar alerts, allowing for the general performance of the SOC in linking similar threats and actualities.

Coordinating workflows and leveraging LLMs for contextual alert analysis

The orchestrator must coordinate all aspects concerning properly sharing data within the system. It relays processed alerts or embedding to the relevant parts of the solution, like LLM or SOAR. Implemented using frameworks, the orchestrator tracks the components' input needs and ensures that each works well. Due to the centralization of management, the orchestrator eliminates idle time and increases system availability and expandability. The primary component of this structure is the large language model. The LLM analyses the enhanced alerts, using its more sophisticated NLP algorithms to read the text of each alert and determine whether it is a legitimate positive or a false positive alert. The LLM refines alert classification through contextual information, historical data, and patterns to minimize false alarms and guarantee that SOC analysts will only be aware of real threats. This step dramatically decreases the amount of work required by an operator, making handling an incident faster and more efficient.

Automating incident response with SOAR and data protection

The LLM sends accurate positive alerts to the SOAR platform where incidents are progressing. The playbooks available on the SOAR platform help predefine an incident response tactic and execute it on autopilot. These playbooks can stamp systems known to be infected, inform the concerned members of an organization or begin forensic examinations. Executing such mundane tasks, the SOAR platform accelerates work or shortens response time and enables SOC analysts to perform other tasks to manage existing threats. Data protection is also a component of the plan blueprint. To facilitate privacy and simplicity of the components, all the data flow transactions between the elements are encrypted through Transport Layer Security/Secure Socket Layer. This encryption aids in protecting identity data from data leakages to third parties and is crucial to enabling businesses to meet their legal obligations as provided under the law. It is also periodically secured through security audits and compliance checks to ensure the system is prepared well enough for more recent threats.

Tools And Technologies Used

The tools and technologies sub-section provides more information about the systems and various platforms used in architecture. These technologies ensure optimal data throughput, higher data quality, and secure coupling of all the components.

SIEM platforms and pre-processing

The SIEM system is the core of the architecture as it contains real-time alerts and logs, which are the basis of the architecture. These systems have a flexible interface and rule set that can customize alert generation systems using the dashboards and rule engines of an organization's security program [7]. With the proposed architecture, it is impossible to have the problem of the

discontinuity or unreliability of the event monitoring carried out by the SIEM. The data pre-processing step involves using tools for extracting, transforming, and loading. These tools are created to process large amounts of raw data, which means they are pre-processed to prepare for analysis. Elements including auto-preprocessing, metadata extraction, and eliminating data redundancy guarantee that the fed data is high quality. These tools significantly contribute to the most effective further analysis of raw security logs since they give accurate results downstream.

Semantic search and workflow management via orchestration frameworks

Semantic embeddings are computed using ready-made models. These embeddings preserve the contextual meaning of the alerts to enable classification and analysis. These embeddings are saved and used in the vector database, providing a fully manageable and efficient search experience. Embeddings and vector databases make performing the semantic query and its correlations easier. Therefore, related threats are easily picked. The orchestrator, which has also been developed, employs programming frameworks to facilitate element communication. It orchestrates and fires out processed alerts and embeddings to their target consumption queues, whether LLM or SOAR system. Thus, the orchestrator minimizes response time and increases the system's reliability in terms of scalability and fault tolerance of the functions.

Advanced NLP with LLMs, automating workflows, and ensuring compliance

A large language model (LLM) is utilized to classify alerts. With enhanced natural language processing (NLP) capabilities, the LLM enriches and analyses data to classify alerts. This integration enables the system to consider specific contextual data effectively, improving functionality and results without significantly burdening SOC analysts. Many SOAR systems define incident response scenarios, executing predefined scripts for specific tasks during an incident, such as system isolation, stakeholder notification, and forensic analysis. SOAR platforms enhance SOC efficiency by automating repetitive tasks and streamlining security operations. Compliance is crucial for securing client-server data exchanges, which rely on end-to-end encryption protocols. Routine security assessments and compliance checks safeguard critical data and ensure that organizations meet regulatory requirements. These measures help protect systems from cyber threats while aligning organizational policies with industry standards.

Proposed Method

The process described in this section is ordinary and step-by-step, but key procedural actions are considered correct in performing security operations. Semantically embedding the data, applying different preprocessing techniques, and selecting improved algorithms would guarantee the creation of a better threat detection system. Besides, its functioning in closed feedback cycles permanently increases the ability to learn new threats from the constantly evolving world.

Enhancing security data processing and automated alert classification

The methodology involves preprocessing raw alerts produced by the SIEM system. These activities include converting the format, eliminating duplicate and noisy data, and including contextual meta-alerts. As a result of the described procedures, an order of magnitude greater and richer set of embedding inputs makes it possible to assess the outcome adequately. The enriched alerts are represented in semantic embeddings, which provide context to each alert and establish its relationship with others. Such embeddings are stored in a vector database and can be easily accessed during semantic search. This approach increases the extent to which the system can relate associated alerts, improving threat detection and categorization. The enriched alerts are then passed to the LLM, where contextual information is processed to flag them as true or false positives. Automation of true positives prioritizes the severity of cases, and the SOAR platform executes response workflows. This approach reduces the burden on SOC analysts and ensures that incident response is efficient and effective.

Modifications through feedback loops

Classified alerts are stored in the system and used as feedback for more new alerts. This makes the system better than the initial one because it updates it every time it has improved its capability to deal with threats.

System Deployment And Implementation

The implementation phase transforms the concept inherent in the methodology into a working system, integrating and utilizing numerous technologies and processes to produce the

intended results. This section defines how the above security enhancement tools, specific configurations, and workflows are implemented sequentially to meet performance and scalability needs. That way, the implementation guarantees a high-performance security operations pipeline that can adapt to dynamic threat environments on industry-leading frameworks and automation platforms.

Optimizing SIEM configuration and pre-processing for scalable analysis

The SIEM system establishes a connection between information through APIs or sustainable event streams to reach the data pre-processing module. Through integration, this system maintains steady security event processing for continuous event supply to analysts. Specific programs run the pre-processing logic through multiple data normalization enrichment and deduplication steps. The data processing protocol combines various methods that guarantee high-quality information output, leading to more precise and higher-performance results. The embedding model builds semantic representations based on its pre-training. An efficient adjacency list graph system has been created to store embeddings and their ID documentation for quick scale-up semantic search capabilities.

Deploying orchestration and AI-driven security automation

The workflow management system and the automation framework are the foundation for building the orchestrator, which directs multiple procedural executions. This component places security event information and embeddings into processing pipelines, which ensures integrated execution functions. An LLM model operates in the environment because APIs link it to the orchestrator for instant alert classification. The integrated system operates better for threat detection through dynamic security event analysis. The SOAR platform performs two main tasks: processing positive alert validation and running predefined incident response templates. Security operation center personnel can dedicate their efforts to high-priority threats because this automated system decreases response time and minimizes human intervention. The system meets security compliance through TLS/SSL encryption technology, and security audit functionalities were added during development. The implemented security measures enable organizations to meet data protection regulatory requirements by stopping unauthorized data release and keeping systems safe.

Challenges And Limitations

To achieve the most effective solutions when organizing a robust and highly scalable security architecture, some problems must be solved, and everyone must be somewhat effective. Stretching from one's ability to grapple with the chaotic mass of big and diversified information flows to the issues related to applying such cutting-edge tools and applications as LLMs, these issues depict the nuances of the interdependence between innovative solutions and business functionality. Furthermore, it discusses the emerging problems anticipated to slow the system implementation process, the impacts to be realized and the steps to strengthen system viability.

Data volume and complexity

Large scale and data heterogeneity are two significant factors that prevent this architecture from being set. Modern systems are flooded with systems by massive volumes of diverse and dissimilar data from sources like users' activities, the network, or system performance. This data is not only characterised by a considerable amount but also by a wide variety: they should be able to accommodate various formats, protocols and ways of data analysis. Therefore, the constant ingestion of data streams increases the computational demand as the system must concurrently analyse the feeds without delays and throughput decline.

False positives in classification

With the integration of LLMs, false positives in threat classification are minimal, but the problem is not eliminated. Cases of misclassification are not yet rare, mainly due to the reliability of datasets used for model training and specific tests that might be beyond the model's perceivable range. The problem with false positives is that they tremendously weaken the system's stability. Since normal activities are considered potential threats, SOC analysts must take action to validate the alerts. This extra load can take the analysts away from the actual threats and may even introduce inefficiencies and cause delays in detecting important events/Incidents. Fixing this problem entails incremental changes to the LLM, more significant training sets, and additional methods to verify the

output of the LLM. A researcher employs a False Positive Rate (FPR) during evaluation to identify the frequency of wrong threat alarms in digital security detection processes. The FPR represents the proportion the system incorrectly labels as potential threats among authentic non-threat events. Mathematically, it is represented as:

$$FPR = \frac{FP}{(FP + TN)} \quad (1)$$

The detection system counts erroneous threat detections as FP False Positives among all monitored everyday activities while classifying legitimate benign events as TP True Negatives. Excessive misclassification, indicated by a high FPR number, causes Security Operations Centers to experience unnecessary alert fatigue. The FPR of LLM-driven SIEM systems can be decreased by improving model training, context incorporation, and threshold refinement, which leads to increased system reliability and efficiency.

Integration overhead

Combining architecture with the number of tools and platforms in a contemporary organizational environment is challenging. The tools or platforms, by design, have different interfaces, data models, and ways of working, making the setup and management mechanical. Most integration maps involve melding old structures into new ones with compatibility problems, data mapping problems, synchronization difficulties or both. Moreover, integration needs to be preserved throughout the architecture's life cycle, which is a very nontrivial task, especially in the presence of upgrades or modifications.

Security vulnerabilities

However, all this sophistication of the architecture does not help guard the system against threats and risks. Such adversarial techniques may include data preprocessing pipelines, training algorithms, or other libraries intrinsic to the system and can be abused. Such subtle attacks are possible; the dataset labelling the model can be poisoned through techniques like poisoning attacks, and the attacker can also use techniques like evasion to evade the defenses. This means that continuous monitoring and updating are required to reduce these risks. We can also include a more detailed defense to enhance the effectiveness of the security measure that will have been deployed by deploying Adversarial training & Penetration testing. However, these measures add to the architecture's operational overhead and resource needs.

Cost considerations

Using such modern solutions as LLMs and vector databases is expensive. However, the costs of training and deploying these systems, together with the costs of achieving high performance in terms of both computing infrastructure and personnel, can be expensive. However, the level of specialization required to manage and enhance the architecture forms part of the overall operating expenses. Managers need to ensure that these technologies bring more benefits than their costs can offset or revise their strategies for using them. Managers need to calculate the long-term return on investment, which is required while searching for a way to spend less, for instance, using cloud services or selecting more effective algorithms.

RESULTS AND DISCUSSION

This paper establishes the importance of embedding LLMs in SOC operations, and the findings are revealed in the subsequent sections. Consequently, when implementing the alert classification, the proposed system fixes gaps such as false positives and analyst workload and enriches contextual understanding. This part provides an overview of the significant findings obtained at the end of the work, accuracy and time improvements and the overall comparison of Artificial Intelligence and Traditional approaches for implementing the solution. Furthermore, the solution's cost efficiency is discovered, and its usability in contemporary cybersecurity operations is reinforced.

Key Results

Low false negative analysis

The implemented system for false positive carcinoma detection has an accuracy of 95.5%, significantly better than traditional techniques. Miscreants have even elaborated deceptive schemes

over time to generate pointless alarms within SOCs that deter analysts from authentic threats. The large language model implemented by the system makes it possible to classify and recognize the alerts accurately. This enhancement will also enable SOC teams to address significant events and reduce cases of slip-through threats.

Reduction in analyst workload

This system also realizes another significant benefit of cutting analyst work time to half through automation and alert consolidation. SOC analysts, therefore, often spend a lot of time scouring through a flood of inconsequential or unrelated alarms. Most of the triage is performed by the LLMs, which means that analysts can attend to complex cases promptly. Not only do they save time and increase SOC's performance, but they also reduce the amount of the analyzing workloads, which may lead to stress in SOC teams.

The Analysis of the Comparative Performance

Rule-based systems vs. ai classification

Conventional validated paper systems employ rigid deterministic rationality with constraints and explicit governing knowledge. These are easy to implement but come with low statistical false positive levels and poor versatility. They are generally insufficiently adapted to account for nonlinear threats or evolving threat vectors. On the other hand, the AI-based classification technique works with sophisticated machine learning algorithms to adjust for newer and emerging threats. The integration of LLMs extends this flexibility by offering great contextual insight and better classification performance, making the system far more robust and consummate than rule-based systems.

Integrated LLM compared to conventional ML methods

Although conventional ML enabled more efficient alert classification than rule-based methods, it is associated with problems such as a limited ability to use new comprehensive training data and the inability to consider the necessary context. However, LLMs are superior in structurally unamorous information, such as security logs and analyst notes, where their semantic analysis is unique.

Table 1. Comparative analysis of LLM integration

Feature	Rule-based systems	AI-driven classification	Traditional ML approaches	LLM integration
Accuracy	Moderate (65-75%)	High (90-95%)	Moderate (80-85%)	Very high (95.5%)
Scalability	Limited	High	Moderate	High
False positive rate	High	Moderate	Moderate	Low
Contextual understanding	Minimal	Advanced	Moderate	Very advanced
Adaptability to new threats	Limited	High	Moderate	High
Ease of deployment	High	Moderate	High	Moderate
Resource requirements	Low	Moderate	Low	High
Analysis of unstructured data	No	Limited	Moderate	Extensive

As illustrated in Table 1, requirements integration to LLMs is fast relative to contextual understanding, improvement in scalability and the resultant enhancement in the accuracy levels categorically underpinning the assertion that integration of LLM implies change in SOC processes at a transformational level.

Cost Efficiency Metrics

Reduction in operation costs

The research showed that operational cost divisions significantly impacted security measures. Using powML with the LLM as an enabled system enables organizations to reduce operational expenses by 30% by eliminating repetitive tasks, including alert sorting and incident investigation. Organizations obtain the most economical way to enhance security via this system, which achieves equal efficiency and effectiveness from reduced costs.

Return on investment (ROI)

Though deploying the LLMs and the supporting infrastructure involves relatively higher initial costs, the system shows great potential for high ROI. The savings from the need for physical effort and the balance of the expenses against potential losses that result from security breaches are

invaluable. For instance, the time and cost saved by cutting the number of analysts by half and the probable financial loss avoided because of threats that have gone unnoticed provide powerful arguments for using this technology.

Improved efficiency in resource allocation

The system makes basic processes run automatically and allows analysts to concentrate on critical threats. It efficiently distributes human and computational power, enhancing long-term cost-value and better organizational productivity—refortifying its proposition.

CONCLUSION

The embracement of Large Language Models (LLMs) to improve cybersecurity work schemes has revolutionized alerts' general filtering and classification. In this study, it can be explained how the proposed framework 'SECiSADH' enhances the sustainability of SOC with 99% promising results in false identification compared to other existing research with an accuracy of 95.5%. The architecture also demonstrates the principles of scalability and real-time work as masses of data are taken and analyzed. The system thus reduces analysts' workload by 50%, allowing SOC teams to filter real threats and improve the organization's security.

The framework's resilience in dynamic threat settings will be further strengthened by the ongoing integration of cutting-edge AI techniques like Reinforcement Learning with Human Feedback (RLHF) and external threat intelligence feeds. Furthermore, the architecture can solve security issues in decentralized and resource-constrained contexts by expanding the application of this approach to new areas, including cloud-native infrastructures and the Internet of Things. Adaptive and intelligent frameworks are improvements and crucial developments for the next generation of SOC operation as cyber threats increase in complexity and volume.

AUTHOR CONTRIBUTIONS

Author of this article played an important role in the process of method conceptualization, formal analysis, and article writing.

CONFLICT OF INTEREST

The author declare that have no conflicts of interest.

REFERENCES

Ali, G., Shah, S., & Elaffendi, M. (2025). Enhancing cybersecurity incident response: AI-driven optimization for strengthened advanced persistent threat detection. *Results in Engineering,* 25, 104078. https://doi.org/10.1016/j.rineng.2025.104078

Alwarafy, A., Al-Thelaya, K. A., Abdallah, M., Schneider, J., & Hamdi, M. (2020). A survey on security and privacy issues in edge-computing-assisted Internet of Things. *IEEE Internet of Things Journal*, 8(6), 4004-4022. https://doi.org/10.1109/JIOT.2020.3015432

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., & Amodei, D. (2020). Language models are few-shot learners. *arXiv preprint arXiv: 2005.14165*. https://doi.org/10.48550/arXiv.2005.14165

Ferrag, M. A., Maglaras, L. A., Moschoyiannis, S., & Janicke, H. (2020). Deep learning for cyber security intrusion detection: Approaches, datasets, and comparative study. *Journal of Information Security and Applications*, 50, 102419. https://doi.org/10.1016/j.jisa.2019.102419

Ferrag, M. A., Ndhlovu, M., Tihanyi, N., Cordeiro, L. C., Debbah, M., Lestable, T., & Thandi, N. S. (2024). Revolutionising cyber threat detection with large language models: A privacy-preserving BERT-based lightweight model for IoT/IIoT devices. *arXiv preprint arXiv: 2306.14263*. https://doi.org/10.48550/arXiv.2306.14263

González-Granadillo, G., González-Zarzosa, S., & Diaz, R. (2021). Security information and event management (SIEM): analysis, trends, and usage in critical infrastructures. Sensors, 21(14), 4759. https://doi.org/10.3390/s21144759

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Aaron, C., & Bengio, Y. (2020). Generative adversarial nets. *Communications of the ACM,* 63(11), 139-144. https://doi.org/10.1145/3422622

Hassanin, M., & Moustafa, N. (2024). A comprehensive overview of large language models (LLMs) for cyber defenses: Opportunities and directions. *arXiv preprint arXiv: 2405.14487*. https://doi.org/10.48550/arXiv.2405.14487

Hyun, T. T., Nguyen, T. D., & Tan, H. (2019). A survey on security and privacy issues of blockchain technology. *2019 International Conference on System Science and Engineering (ICSSE)*. https://doi.org/10.1109/ICSSE.2019.8823094

Kshetri, N., Pandey, P. S., & Ahmed, M. (2025). *Blockchain technology for cyber defense, cybersecurity, and countermeasures*. CRC Press. https://doi.org/10.1201/9781003449515

Liang, X., Zhao, J., Shetty, S., Liu, J., & Li, D. (2017). Integrating blockchain for data sharing and collaboration in mobile healthcare applications. *2017 IEEE 28th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC),* 46-57. https://doi.org/10.1109/PIMRC.2017.8292361

Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *arXiv preprint arXiv: 1705.07874*. https://doi.org/10.48550/arXiv.1705.07874

Moustafa, N., & Slay, J. (2016). The evaluation of network anomaly detection systems: Statistical analysis of the UNSW-NB15 dataset and the comparison with the KDD99 dataset. *Information Security Journal: A Global Perspective*, 25(1-3), 18-31. https://doi.org/10.1080/19393555.2015.1125974

Pathak, M., Mishra, K. N., & Singh, S. P. (2024). Securing data and preserving privacy in cloud IoT-based technologies an analysis of assessing threats and developing effective safeguard. *Artificial Intelligence Review*, 57(269), 1-46. https://doi.org/10.1007/s10462-024-10908-x

Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). Improving language understanding by generative pre-training. *OpenAI preprint*. 1-12. Retrieved from https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf

Reddy, V. S. S., & Reddy, N. (2024). AI-powered language models enhance natural language understanding and generation. International Journal of Artificial Intelligence & Machine Learning (IJAIML), 3(02), 101-115. https://lib-index.com/index.php/IJAIML/article/view/IJAIML_03_02_008

Salah, K., Rehman, M. H. U., Nizamuddin, N., & Al-Fuqaha, A. (2019). Blockchain for AI: Review and open research challenges. *IEEE Access*, 7, 10127-10149. https://doi.org/10.1109/ACCESS.2018.2890507

Shaukat, K., Luo, S., Varadharajan, V., Hameed, I. A., & Xu, M. (2020). A survey on machine learning techniques for cyber security in the last decade. *IEEE Access*, 8, 219308-219323. http://dx.doi.org/10.1109/ACCESS.2020.3041951

Sun, P. (2020). Security and privacy protection in cloud computing: Discussions and challenges. *Journal of Network and Computer Applications*, 160, 102642. https://doi.org/10.1016/j.jnca.2020.102642

Yadav, T., & Rao, A. M. (2015). Technical aspects of cyber kill chain. *Proceedings of the International Conference on Advances in Computing, Communications and Informatics (ICACCI)*, 2354-2360. http://dx.doi.org/10.1007/978-3-319-22915-7_40