International Journal of Computer Science and Information Technology Research (IJCSITR)
2025, Vol. 6, No. 1, January - February, pp. 72-82
Journal ID: 9471-1297
DOI: https://doi.org/10.63530/IJCSITR_2025_06_01_008
website: www.ijcsitr.com

Autonomous Cyber Defense: LLM-Powered Incident Response with LangChain and SOAR Integration

Sandhya Guduru
Masters in Information Systems Security,
Software Engineer - Technical Lead, USA.

Abstract

The increasing sophistication of cyber threats necessitates the adoption of advanced, autonomous defense mechanisms. Large Language Models (LLMs) have emerged as a powerful tool for automating cybersecurity workflows, enabling intelligent incident response. This paper explores integrating LLM-powered incident response using LangChain, a framework that enhances natural language processing capabilities, and Security Orchestration, Automation, and Response (SOAR) platforms like Tines for automated containment workflows. The proposed system leverages MITRE ATT&CK playbooks to train LLMs, ensuring contextual decision-making and threat mitigation. Furthermore, probabilistic graphical models (PGMs) validate LLM-driven decisions, enhancing reliability and reducing false positives. This approach minimizes response time and enhances cybersecurity resilience by automating threat detection, triage, and containment. The findings underscore the transformative potential of AI-driven cyber defense, offering a scalable and efficient solution for mitigating modern cyber threats.

Keywords

Autonomous Cyber Defense, Large Language Models (LLMs), LangChain, Security Orchestration Automation and Response (SOAR), MITRE ATT&CK, Probabilistic Graphical Models (PGMs), AI-Driven Incident Response, Cybersecurity Automation.

How to Cite: Sandhya Guduru. (2025). Autonomous Cyber Defense: LLM-Powered Incident Response with LangChain and SOAR Integration. *International Journal of Computer Science and Information Technology Research (IJCSITR)*, 6(1), 72-82.

DOI: https://doi.org/10.63530/IJCSITR_2025_06_01_008

Article ID: IJCSITR_2025_06_01_008

Introduction

In today's digital landscape, organizations face an escalating number of sophisticated and persistent cyber threats. Traditional cybersecurity measures often struggle to keep pace with the rapid evolution of attack vectors, resulting in increased vulnerabilities. According to recent statistics, it takes an average of 258 days for security teams to identify and contain a data breach, highlighting the inefficiency of current incident response strategies.

Integrating Artificial Intelligence (AI) into cybersecurity has emerged as a transformative approach to enhance threat detection and response capabilities. AI technologies, such as machine learning and natural language processing, enable the analysis of vast datasets to identify patterns and anomalies indicative of malicious activity. For instance, AI-driven systems have been successfully implemented to automate incident response, reducing the burden on human analysts and improving reaction times [1].

Large Language Models (LLMs) represent a significant advancement in AI, offering the ability to process and generate human-like text based on extensive training data. In cybersecurity, LLMs can be leveraged to interpret complex threat intelligence, generate comprehensive reports, and even suggest remediation steps, thereby streamlining the incident response process. Recent case studies have demonstrated the effectiveness of LLMs in enhancing cyber defense mechanisms, showcasing their potential to revolutionize traditional security operations [2].

This paper explores the integration of LLMs into cybersecurity incident response

frameworks, focusing on their role in automating threat detection, analysis, and mitigation. By examining existing approaches, recent advancements, and the challenges associated with AI-driven security solutions, we aim to provide insights into developing more resilient and adaptive cyber defense strategies.

Literature Review

Integrating Artificial Intelligence (AI) into cybersecurity has revolutionized incident response mechanisms, enabling automated detection and mitigation of threats. AI-driven systems have demonstrated the capability to reduce the mean time to detect (MTTD) by 35% and the mean time to respond (MTTR) by 42%, enhancing overall cybersecurity resilience [3].

Large Language Models (LLMs), such as GPT-4, have emerged as pivotal tools in cybersecurity. Their ability to process and generate human-like text facilitates advanced threat detection and response strategies. Recent studies have systematically investigated the application of LLMs within the field, covering over 180 academic papers since 2023 providing a comprehensive overview of their current state, challenges, and future directions [4].

The LangChain framework has been introduced to enhance the development of LLMs-powered applications. LangChain offers a versatile and modular approach, simplifying complex stages of the application lifecycle, such as development, productionization, and deployment, thereby facilitating the creation of scalable and contextually aware applications [5].

Integrating AI into Security Orchestration, Automation, and Response (SOAR) platforms has further optimized incident response processes. AI-driven SOAR systems can automate tasks, enhance threat detection, and improve response accuracy, addressing challenges like adversarial attacks and the need for high-quality data [6].

Probabilistic Graphical Models (PGMs) have been employed to validate decisions within AI-driven incident response systems. These models provide a structured approach to reasoning under uncertainty, enhancing the accuracy and reliability of automated responses [7].

Despite these advancements, challenges persist in integrating AI into cybersecurity. Data quality, ethical considerations, and the potential for adversarial attacks necessitate ongoing research and development to fully realize the potential of AI-driven incident response systems [8].

In summary, the convergence of AI technologies, including LLMs, frameworks like LangChain, and AI-enhanced SOAR platforms, has significantly advanced cybersecurity incident response. Continued exploration and integration of these technologies are essential to effectively address emerging cyber threats.

Problem Statement

In today's digital landscape, organizations face an escalating number of increasingly sophisticated and diverse cyber threats. Traditional incident response (IR) methods often struggle to keep pace with these evolving threats, leading to prolonged detection and remediation times. A study by IBM revealed that only 26% of organizations have an incident response plan that is consistently applied, highlighting a significant gap in preparedness [9].

The complexity of managing multiple stakeholders further complicates incident response efforts. Diverse expectations and communication requirements among IT teams, business units, customers, and external partners can lead to misunderstandings and delays. A survey by Forrester indicated that 64% of IT leaders reported differing expectations regarding incident resolution times among stakeholders, while 57% noted varying communication needs [9].

Resource limitations and budget constraints exacerbate these challenges. Effective incident management demands specialized skills and tools, which can be costly to acquire and maintain. The ITIL Foundation reported that 70% of IT managers identified budget constraints as a significant barrier to effective incident management, with 55% citing a lack of resources as a significant challenge [9].

Furthermore, the rapid adoption of emerging technologies such as cloud computing, artificial intelligence, and the Internet of Things introduces new complexities. These technologies necessitate specialized skills and processes, which can be challenging to implement and manage effectively. A survey by 451 Research found that 60% of IT managers viewed emerging technologies as a significant challenge to incident management, with 55% indicating a need for new skills and training [9].

In light of these challenges, there is a critical need for innovative solutions to enhance incident response efficiency and effectiveness. Integrating advanced technologies, such as Large Language Models (LLMs) powered by AI, with frameworks like LangChain and Security Orchestration, Automation, and Response (SOAR) platforms offers a promising

avenue to address these limitations. By automating routine tasks, improving communication among stakeholders, and providing intelligent decision support, these integrated solutions can revolutionize incident response in the face of an ever-evolving cyber threat landscape.

Existing Approaches and Limitations

Traditional cybersecurity incident response relies heavily on rule-based systems, signature detection, and human analysts to identify and mitigate threats. Security Information and Event Management (SIEM) tools and Security Orchestration, Automation, and Response (SOAR) platforms aggregate logs and automate certain workflows but still require predefined rules and manual oversight. While effective against known threats, these systems struggle with novel attack techniques that do not match existing signatures.

Behavior-based anomaly detection employs machine learning to identify deviations from regular network activity. However, such systems often generate many false positives, requiring human analysts to validate and respond to alerts. The reliance on security teams for threat triage, investigation, and response slows down remediation efforts, increasing the risk of cyberattack damage. Moreover, adversaries continuously evolve their tactics, techniques, and procedures (TTPs), rendering static detection models obsolete.

Current approaches also lack contextual awareness when correlating attack signals across multiple sources. Threat intelligence feeds improve detection capabilities but do not provide automated reasoning to assess attack severity or recommend the best action. Consequently, security teams experience alert fatigue and critical threats may go unnoticed due to the overwhelming volume of data [10].

Automation integration in cybersecurity has improved efficiency, but the limitations of traditional methods leave organizations vulnerable to sophisticated threats. These gaps highlight the need for a more adaptive, intelligent, and autonomous incident response system to dynamically analyze and respond to threats in real-time.

Advancements in AI for Cyber Defense

Recent artificial intelligence (AI) advancements have significantly enhanced cybersecurity measures by enabling automated threat detection and response. Machine learning models can analyze vast amounts of data to identify patterns and anomalies indicative of cyber

threats, facilitating real-time monitoring and proactive defense mechanisms [7].

Large Language Models (LLMs) have revolutionized cyber defense by providing contextual analysis and intelligent decision-making capabilities. By integrating machine learning-based anomaly detection with explainable AI, LLMs can assist in intrusion detection and offer understandable explanations for their outputs, thereby enhancing the interpretability and trustworthiness of AI-driven security systems [11].

Natural Language Processing (NLP) techniques empower LLMs to automate incident triage by summarizing security logs, correlating attack signals, and prioritizing high-risk alerts. This reduces the burden on security teams and improves response efficiency. Additionally, AI-driven chatbots and virtual assistants enhance cybersecurity operations by guiding analysts through remediation workflows and automating threat-hunting processes.

Integration with Security Orchestration, Automation, and Response (SOAR) platforms amplifies AI's impact by automating responses to threats, such as blocking malicious traffic and isolating compromised devices, thereby saving time and minimizing data breaches [12].

Despite these advancements, AI-driven cybersecurity solutions face challenges. Adversaries also leverage AI to craft sophisticated attacks, such as AI-generated phishing emails and deepfake-based social engineering. Moreover, AI models require continuous training on up-to-date threat intelligence to remain effective. Biases in training data can lead to misclassification of threats, and adversarial machine learning techniques can manipulate AI-driven security systems [6].

Figure 1: Traditional vs. AI-Driven Incident Response

Aspect	Traditional Incident Response	AI-Driven Incident Response
Threat Detection	Rule-based detection, high false positives	AI-driven anomaly detection, reduced false positives
Response Speed	Manual, slow response to threats	Real-time automated containment
Decision Making	Human analysts decide based on predefined rules	Machine learning models analyze and respond
Scalability	Limited, depends on workforce availability	Highly scalable, adapts to workload
Accuracy	Prone to errors due to human fatigue	Enhanced precision with continuous learning
Automation Level	Minimal automation, high manual intervention	Fully automated, human intervention minimal
Adaptability	Static rules, requires frequent updates	Dynamic learning, adapts to new threats

Proposed Solution

Figure 2: Automated Threat Containment Workflow using LLMs and SOAR Integration.

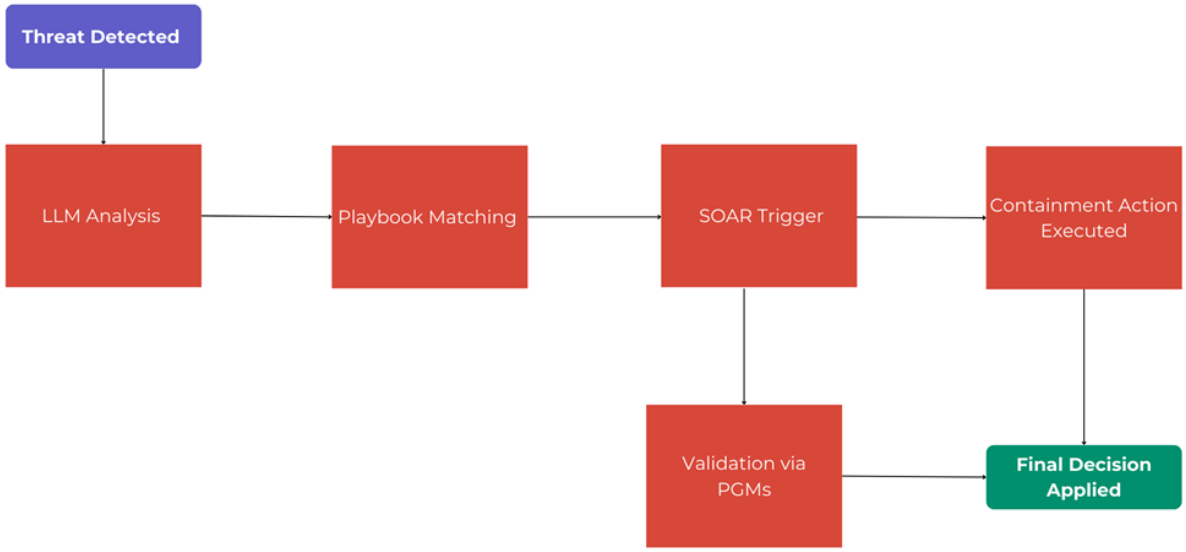

To address the challenges identified in the problem statement, we propose an integrated approach that combines Large Language Models (LLMs) with Security Orchestration, Automation, and Response (SOAR) platforms. This integration aims to enhance the efficiency and effectiveness of incident response processes within Security Operations Centers (SOCs).

Leveraging Large Language Models (LLMs) in Incident Response

LLMs, such as GPT-4, have demonstrated significant natural language understanding and generation capabilities. In the context of cybersecurity, these models can be trained on frameworks like the MITRE ATT&CK playbooks to generate contextually relevant and actionable responses to security incidents. By analyzing vast amounts of data, LLMs can assist in identifying patterns and anomalies indicative of cyber threats, thereby improving detection accuracy. Furthermore, LLMs can automate the generation of incident reports and recommend remediation steps, reducing the cognitive load on human analysts and expediting response times.

Integration with LangChain for Enhanced Natural Language Processing

LangChain is a framework designed to facilitate the development of applications powered by language models. By integrating LLMs with LangChain, SOCs can create more dynamic and responsive incident response systems. This integration enables the parsing and interpreting complex security alerts and logs, allowing for more accurate and context-aware responses. Additionally, LangChain's capabilities can be leveraged to develop conversational interfaces, enabling analysts to interact with the system using natural language queries, thus streamlining the investigative process.

Synergy with SOAR Platforms

SOAR platforms are designed to improve the efficiency of SOCs by automating routine tasks, orchestrating workflows, and facilitating coordinated responses to security incidents. Integrating LLMs into SOAR platforms enhances these capabilities by introducing advanced decision-making processes based on comprehensive data analysis. For instance, LLMs can prioritize alerts based on their potential impact, suggest optimal response strategies, and adapt to evolving threats in real-time. This integration accelerates response times and ensures that a deep understanding of the current threat landscape informs actions taken [13].

Validation through Probabilistic Graphical Models (PGMs)

To ensure the reliability and accuracy of decisions made by the integrated system, we propose the use of Probabilistic Graphical Models (PGMs). PGMs provide a framework for modeling the probabilistic relationships among variables in complex systems. By incorporating PGMs, the system can assess the confidence levels of its predictions and recommendations, thereby providing analysts with insights into the certainty of suggested actions. This validation step is crucial in maintaining trust in automated systems and ensuring that critical decisions are made with a clear understanding of associated uncertainties.

Figure 3: AI-Powered Incident Response vs. Traditional Methods

Feature	Traditional Incident Response	SOAR-Based Response	LLM-Powered (LangChain + SOAR)
Automation Level	Manual, rule-based scripts	Automated playbooks	Dynamic, AI-generated responses
Threat Detection Speed	Slow due to manual investigation	Faster but dependent on pre-set rules	Real-time, adaptive detection
Decision Making	Human analysts interpret data	Automated but rigid decision-making	AI-driven contextual decision-making
Natural Language Processing (NLP)	Not used	Limited NLP for log analysis	Advanced NLP for threat analysis
Adaptability to New Threats	Limited, requires frequent updates	Can adapt to some extent	Highly adaptable with self-learning models
Integration with SIEM	Manual correlation	Predefined integrations	AI-driven log analysis and automated response

The proposed integration of LLMs with LangChain and SOAR platforms offers a comprehensive solution to the challenges faced by modern SOCs. By harnessing the advanced capabilities of LLMs in natural language processing and data analysis, combined with the automation and orchestration strengths of SOAR platforms, this approach aims to enhance the speed, accuracy, and effectiveness of incident response processes. Incorporating PGMs further ensures that decisions are validated and trustworthy, ultimately contributing to a more resilient cybersecurity posture.

Conclusion

The integration of Large Language Models (LLMs) into cybersecurity frameworks represents a transformative step in automating threat detection and incident response. By leveraging frameworks like LangChain, organizations can enhance the capabilities of LLMs, enabling more sophisticated analysis and response strategies. This approach facilitates the development of AI-driven Security Operations Centers (SOCs) that proactively simulate, detect and mitigate cyber threats.

Furthermore, the incorporation of Security Orchestration, Automation, and Response

(SOAR) platforms amplifies incident management efficiency. AI-powered workflows seamlessly integrate into existing security infrastructures, providing real-time remediation suggestions and creating an end-to-end automation ecosystem.

Despite these advancements, challenges remain, particularly in ensuring the security and reliability of AI-driven cybersecurity solutions. Potential vulnerabilities within frameworks like LangChain highlight the need for continuous monitoring and updating of AI models to mitigate risks effectively.

In conclusion, the strategic integration of LLMs, LangChain, and SOAR platforms offers a promising pathway toward more resilient and adaptive cyber defense mechanisms. By embracing these technologies, organizations can enhance their ability to detect, analyze, and respond to cyber threats in an increasingly complex digital landscape.

References

[1] "AI in Cybersecurity: 13 Examples and Use Cases," Perception Point, Nov. 25, 2024. Available: https://perception-point.io/guides/ai-security/ai-in-cybersecurity-examples-use-cases/?.

[2] A. Jamil, "Case Studies: Successful Implementations of AI in Cyber Defense," Umetech.net, Sep. 03, 2024. Available: https://www.umetech.net/blog-posts/successful-implementations-of-ai-in-cyber-defense?.

[3] S. Chahal, "AI-Enhanced Cyber Incident Response and Recovery," International journal of science and research, vol. 12, no. 3, pp. 1795–1801, Mar. 2023, doi: https://doi.org/10.21275/sr231003163025

[4] J. Zhang, "When LLMs Meet Cybersecurity: A Systematic Literature Review," Arxiv.org, 2023. Available: https://arxiv.org/html/2405.03644v1?.

[5] V. Mavroudis, "LangChain," Nov. 2024, doi: https://doi.org/10.20944/preprints202411.0566.v1. Available: https://hal.science/hal-04817573/.

[6] I. Jada and T. O. Mayayise, "The impact of artificial intelligence on organisational cyber security: An outcome of a systematic literature review," Data and Information Management, vol. 8, no. 2, pp. 100063–100063, 2023, doi:

https://doi.org/10.1016/j.dim.2023.100063. Available: https://www.sciencedirect.com/science/article/pii/S2543925123000372

[7] A. H. Salem, S. M. Azzam, O. E. Emam, and A. A. Abohany, "Advancing cybersecurity: a comprehensive review of AI-driven detection techniques," Journal Of Big Data, vol. 11, no. 1, Aug. 2024, doi: https://doi.org/10.1186/s40537-024-00957-y. Available: https://journalofbigdata.springeropen.com/articles/10.1186/s40537-024-00957-y

[8] R. Kelly, "Simplifying Cyber Incident Response Strategies," NormCyber, Nov. 04, 2024. Available: https://www.normcyber.com/blog/navigating-the-complexity-of-cyber-incident-response/?.

[9] admin, "The Limitations of Incident Management: Challenges and Opportunities," CIO Insight Hub, Nov. 07, 2022. Available: https://ciohub.org/post/2022/11/the-limitations-of-incident-management/?

[10] J. Chukwube, "Challenges to Traditional Methods of Cybersecurity Information Gathering - Risk and Resilience Hub," Risk and Resilience Hub, Jan. 26, 2023. Available: https://www.riskandresiliencehub.com/challenges-to-traditional-methods-of-cybersecurity-information-gathering/?.

[11] T. Ali and P. Kostakos, "HuntGPT: Integrating Machine Learning-Based Anomaly Detection and Explainable AI with Large Language Models (LLMs)," arXiv.org, 2023. Available: https://arxiv.org/abs/2309.16021?.

[12] J. Kerwin, "What Is the Role of AI in Cybersecurity?," Excelsior University, Jul. 2024. Available: https://www.excelsior.edu/article/ai-in-cybersecurity/?

[13] G. Sweny, "How AI and LLMs change SOAR and the Security Operations Center (SOC)," Agileblue.com, 2019. Available: https://agileblue.com/how-ai-and-llms-change-soar-and-the-security-operations-center-soc/?