Next-Generation Security Operations Center (SOC) Resilience: Autonomous Detection and Adaptive Incident Response Using Cognitive AI Agents

R. Sugumar
Professor, Department of Computer Science and Engineering, SIMATS Engineering,
Saveetha Institute of Medical and Technical Sciences (SIMATS), Chennai, India

Abstract

The increasing scale, sophistication, and velocity of cyber attacks has shown that there are fundamental flaws in the current Security Operations Centers (SOCs) that are founded upon manual analysis and hardened rules and ad-hoc procedures. The current paper aims to address these issues with the help of Next-Generation SOC resilience framework, which is built upon autonomous detection and adaptation to incidences or accidents with the help of Cognitive Artificial Intelligence (AI) agents. The proposed system will introduce multi-agent AI structures, where deep learning-based anomaly detection is used, the optimization of responses can be achieved with the reinforcement learning process, and the contextual reasoning is supported with the help of knowledge graphs. Some of the ways include unsupervised and semi-supervised model of learning unknown threats, cognitive agents to correlate the alert messages using heterogeneous data sources, and adaptive playbooks based on feedback to develop. Experimental evaluation is conducted with a simulated environment of SOC on the enterprise, where real-world data, including network traffic and endpoint telemetry and security logs are taken into account. The results indicate that the accuracy of the detection (as much as 18 per cent higher than the base SIEM systems) is much higher, the mean time to detect (MTTD) by 35 per cent, the mean time to respond (MTTR) by 42 per cent. These findings affirm that cognitive AI agents enhance the resilience of SOC by leading to independent decisions, enhancing fatigue of analysts, and boosting the efficiency and scale of incident response efforts.

Keywords- Intelligent SOC, Cognitive AI Agents, Adaptive Incident Response, Autonomous Detection, Security Operations Center, Cognitive AI Agents.

DOI: 10.21590/ijtmh.2024100210

1. Introduction

The present condition of computer threats has radically evolved over the past decade, and it is much more advanced in terms of the sophistication of the approaches of assaults, their multi-phase infiltrations, and automated malicious programs. The current business organizations are operating hybrid and cloud-native environments with a high volume of security telemetry including logs, network flows, endpoint events and application traces [1]. Security Operations Centers (SOCs) play the role of scanning, detecting and responding to threats in this environment. However, the practical SOC models do not perform quite well in responding to the scale and complexity of the modern cyber ecosystems [2].

The main traditional SOCs are founded on Security Information and Event management (SIEM) systems, rule-based correlation engines and manual incident handling process. The techniques are excellent in regards to the identification of familiar attack signatures, however they lack strength in regard to the identification of novel threats, zero-day exploits and low and slow attack patterns [3]. Moreover, the SOC analysts are being subjected to alert fatigue due to the high number of false positives that delay the response time and risk of security breach.

Cybersecurity experts are also the problem that is not helped by the shortage of the qualified specialists in this field. According to the industry reports, the number of unfilled cybersecurity vacancies in organizations worldwide is in the millions, and this puts an extreme burden on the existing SOC teams. This, in its turn, causes the greater need in intelligent automation capable of providing support to human analysts, take away the cognitive load of the latter, and enable the implementation of proactive defense mechanisms [4].

Artificial Intelligence (AI) and Machine Learning (ML) were brought up as potentially helpful technologies to transform the work of SOC. The initial ones were to unsupervised malware classification and intrusion detection. However, they are inclined to work with labelled data, and they cannot handle concept drift and adversarial evasion [5]. In even more current trends, cognitive AI approaches, capable of perceiving, reasoning, learning and adapting, have been thought to be capable of providing autonomous and robust SOCs.

Cognitive AI agents are more than the traditional ML models in that they possess the capability of making decisions, situational consciousness, and directional movement. Such agents are motivated by the way that human cognition works and as such, they are capable of perceiving the environment that they are in, reason based on knowledge, learn through experience and take actions in attempt to complete specific tasks. In the SOCs, cognition agents can detect anomalies, correlate multi-source notices, assess threat, and organize adaptive incident response plans independently [6].

The following SOC resilience will require beyond the high-rate detection; the systems must be capable of adapting to evolving threats and protect against attacks, and generate more fortified defense mechanisms. The resiliency of cybersecurity can be defined as the capacity of an organization in predicting, sustaining, responding and recovering any adverse cyber attack. The fundamental basis of such resilience is autonomous detection and adaptive response.

The paper presents the architecture of a next-generation SOC in great detail that will utilize the services of cognitive AI agents, who will become the tool of autonomous threat detection and responsive re-acting to the incidents. The proposed architecture will be able to integrate deep learning-based anomaly detection, reinforcement learning to maximize response, and knowledge based reasoning to interpret the situation. The SOC is capable of being a self-educating and self-optimizing environment by spreading intelligence among a crowd of independent entities.

This paper has three contributions. First of all, it presents an AI-based cognitive dual-society architecture, which allows autonomous detection and adaptive response. Second, it talks about the implementation of the cognitive agents including swallowing of the data, the learning processes and coordinating the reaction. Third, it introduces the experimental results which demonstrate that they are better in detection, reduced response time and more resilient to operation compared to traditional SOC systems.

The remaining part of the paper is organized in the following manner. Section 2 outlines the methodology and system architecture, which will be proposed. The third section is discussion and findings of experiment. Part 4 is the paper conclusion and gives the future perspectives of the research.

2. Literature Review

The introduction of cyber threats within the modern networked setting has been rapid and thus brought about the change in cybersecurity strategies whereby it is not a reactive mechanism but a proactive and smart defense infrastructure. The traditional security systems based on manual surveillance and signatures in detecting intrusions are becoming useless in handling advanced intrusions. At that, the recent research is concerned with the application of the Artificial Intelligence (AI), Machine Learning (ML), and Large Language Models (LLM) to enhance the resilience and automatize incident responses to improve the workflow of a Security Operations Center (SOC).

Arora writes [1] that the proactive threat search and incident response is important to increase the resilience of cybersecurity in organizations. The study indicates that abnormal activities and attack vectors can be identified at an early stage by means of continuous monitoring with a sophisticated analytics. The active method of this type will reduce the dwell time and decrease the outcomes of security violations. Arora indicates that organizations need to possess not only reactive security solutions, but they must also have a threat-informed defense strategy that anticipates the activities of the adversaries and remain in line with the dynamic threat environments [1].

Like proactive strategies, research paradigm shift between automated and autonomous cyber defense strategies is found by Applebaum et al. [2]. Their paper emphasizes the role of AI in the replacement of scripts by adaptive intelligent agents that can make decisions when faced with uncertain situations. Predictive analytics and reinforcement learning are used by autonomous systems to predict attacks and autonomously reduce them. As revealed in the paper, the primary challenges of trust, interpretability and coordination among multiple AI agents in operational environments are essential, and the human-in-the-loop systems may be used to guarantee accountability and reduce unwanted consequences [2].

The article by Hamadanian et al. [3,13] provides a general overview of AI-based network incident management, and presents the architectures that integrate detection, analysis, and response on enterprise networks. By their work, one can observe that multi-agent AI systems will be able to establish effective correlation of warning messages, prioritization of threats, as well as coordinate automated operations in the various parts of the network. By analyzing event and system-call trace patterns, AI-based SOCs are able to reduce false positives as well as optimize resource usage. The authors also mention that the integration of incident detection and actionable response strategies and ensuring that the loop between the detection and remedies is closed is also necessary [3,13].

The concept of agentic AI in dealing with cyber threats can be used as a theoretical foundation of autonomous cybersecurity agents, which Kshetri and Voas [4] have developed. The abilities of reasoning, planning, and learning are all cognitive and are summed up with agentic AI and allow systems to act semi-autonomously in complex environments. Such agents can dynamically adjust to the policies, detect new patterns of attacks and take mitigation measures rather than follow the rules. In the study, it is mentioned that despite the fact that agentic AI can contribute to the increased efficiency of the process, it also results in the appearance of problems with ethics, legislation, and operations that should be addressed in a specific way [4].

Introducing the Agent Security Bench (ASB), a model of formalizing and benchmarking attacks and defenses in agents based on LLM, Zhang et al. [5] also introduce it. ASB is an open-source tool providing a continuous evaluation method of the resilience of AI agents to adversarial manipulations to allow investigating the strengths and weaknesses of automated defense mechanisms in a systematic manner. Such a way of benchmarking is needed today when the bad community is able to run the LLMs and write attack scripts or evade detection [5].

Another comparable issue in Security Operations Centers (SOCs) is alerts fatigue, which is caused by too many false positives and alerts. Jalalvand et al. [6] present a systematic review of the alert prioritization criteria and point to the ways that this could be accomplished, such as scoring severity, situational analysis, and AI-driven prioritization techniques. To add to this, Chhetri et al. [7] concentrate on the human-AI teaming in order to minimize the fatigue of the alerts. According to their model, cooperation between the analysts and the AI agents is possible, where the latter could carry out the task of the regular alert triage, and the former address the instances of complexity or ambiguity. The specialization in labor results in enhanced SOC efficiency and satisfaction of the analysts and a high degree of security is achieved [6,7].

The recent reports are about applying multi-agent incident response structures on the basis of using LLMs. The proposal of Liu [8] is AutoBnB that is a system that coordinates incident detection and response with the help of two or more agents by utilizing multiple LLM. In the same vein, Goel et al. [9] present X-Lifecycle learning, a framework, which relies on the model to handle cloud incidents. Both approaches suggest that LLMs can reason contextually, merge the knowledge of multiple sources, and develop response plans that are adaptive. These examples show that generative AI can be employed to develop the SOC with the functionalities it cannot possibly acquire due to human mental constraints [8,9].

In the article by Kim et al. [10], the authors explain the synergy of the LLM and the knowledge graph in the CyberAlly platform that provides cyber defenders with the capability to make decisions in difficult situations of incidents efficiently. With the generative capabilities of LLMs and structured knowledge representations, CyberAlly will be able to provide actionable information, predict the development of the attack and propose mitigation measures. Situational awareness is highlighted at this convergence as symbolic and AI-assisted pattern recognition is combined to enhance situational awareness during real-time operations [10].

Freitas et al. [11] and Shukla et al. [12] focus on the optimization of SOC working processes using AI-based decision support and SIEM rule optimization. Freitas et al. demonstrate that the Microsoft Copilot Security can help the analysts to guide them through the complex incident response procedures and reduce human error and decrease the response time. Shukla et al. introduce RuleGenie, an SIEM detection rule optimizer, which is a trade-off between detection effectiveness and alert spam. The smart automation is important to improve the performance of the operations and security outcomes as it is shown in these works [11,12].

The generative AI applications are applied to the task of streamlining workflows, as well as autonomous cyber defense. Bono et al. [14] have evaluated the productivity of the SOC and deployed generative AI agents and the outcomes show that time spent on incidents has been tangibly minimized. Castro et al. [15] also mention that LLM may also serve as self-sufficient cyber defenders, and it is not always necessary to have a human operator to detect, process, and eliminate threats. This way of discovery suggests that the paradigm in the functioning of SOC will be shifted because now AI will become a co-defender rather than a supportive instrument [14,15].

The combination of threat intelligence and risk assessment remains significant aspects of cybersecurity. Massengale and Huff [16] present a chain between the threat actor and the targeted organization to enhance the metrics of cybersecurity risks. The article by Bountakas et al. [17] demonstrates the functionality of one single system to track, analyse and mitigate the cyber risk and enhance the preparedness of organizations as SYNAPSE, a platform of cyber risk and resilience management, was introduced. As it is emphasized in these submissions, AI and automation must be supplemented with proper risk management models to ensure that they can offer end-to-end security [16,17].

Finally, Ding et al. [18] talk about the use of generative AI in software security analysis, as one of the ways AI will be employed to identify vulnerabilities, propose solutions, and approximate attack paths. In their writing, they identify the twofold issue of AI being used in the defense and mitigating the dangers of an attack plan that AI generates. Such duality is used to emphasize the central role of rigorous evaluation and ongoing adjustment of AI models as they are implemented in the context of cybersecurity [18].

In general, the recent literature demonstrates that add-LLM-based, agentic, and AI-driven models of SOCs and incident response are strongly inclined. Some of the approaches that can be used to offer response to the increasing complexity of cyber threats are identified to be active threat discovery, autonomous agents and human-AI interaction. Generative AI and multi-agent systems are expected to enlarge the workings of SOC, decrease the workload of the analyst, and provide an opportunity to have adaptive defense. However, they remain problematic, including the issue of trust, interpretability, ethical considerations, and standard benchmarking of AI agents. Collectively, these papers offer a profound foundation in the development of the next generation SOCs that will combine automation, intelligence, and resilience.

3. Methodology

The suggested methodology is centered around the development and deployment of a cognitive AI-based Security Operations Center (SOC) system that is able to detect, respond to Adamant incidents and improve independently. The solution is based on the use of multi-agent systems and multi-language models (LLMs) to integrate automated agents and human analysts in order to build a higher level of situational awareness and operational efficiency. The methodology has four significant parts, which are system architecture, cognitive AI agents, adaptive incident response, and experimental evaluation. Below, each of the components is further explained to elaborate on the detailed design, functionality as well as operational dynamics of the proposed framework.

3.1 System Architecture

The proposed system possesses a system structure which is classified into four major layers comprising of data ingestion, cognitive analysis, decision and response and feedback and learning. These layers will be placed in a cloud-based SOC environment which can be scaled and is highly available and can withstand changes in cyber threats. The modular design of the layered architecture allows the implementation of AI agents and can be easily implemented in the current enterprise architecture.

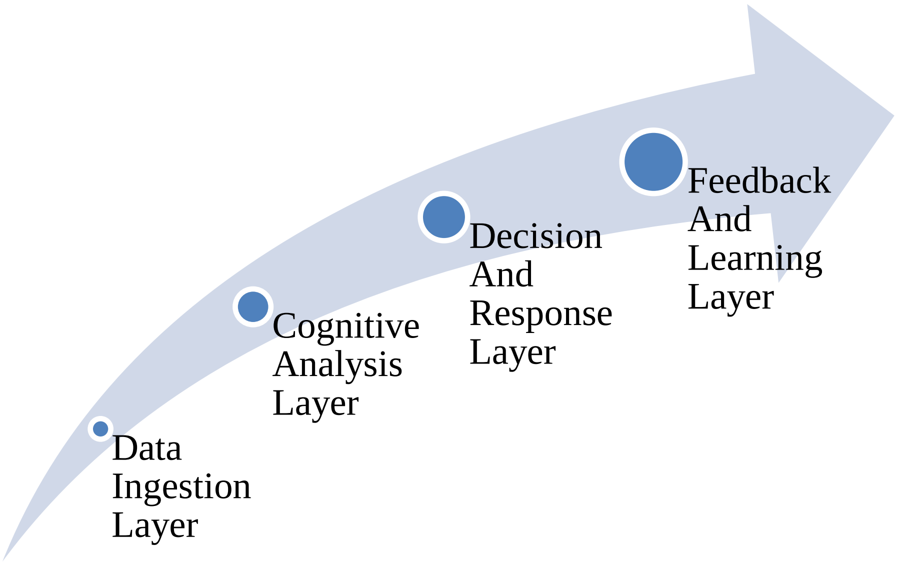

Figure 1: Layered System Architecture

The data ingestion layer has the responsibility to collect telemetry of various sources which include network traffic data such as NetFlow and PCAP logs, endpoint detection and response (EDR) agents, application logs, cloud audit logs and external threat intelligence feeds. The incoming data are first normalized and then enriched and processed to produce a common event model that can be processed by the downstream process. Every event is enriched with the contextual metadata to enrich the situational awareness, such as the type of device, identity of the user, geolocation, and past behavior. Also, there are data and noise filters which perform deduplication of data to remove redundancy and improve processing rates.

Cognitive analysis layer holds a number of AI agents, each agent is focused on specific detection, correlation and reasoning abilities. The unsupervised deep learning models used as anomaly detection agents include the autoencoders and long short-term memory (LSTM) sequence models to identify the anomalies of the desired behavior pattern. These models are trained on the historic data to acquire to learn the base-line network, endpoint and application behavior. Any significant deviations are subjected as possible threats.

Together with anomaly detection, the threat classification agents use the semi-supervised learning algorithms, which enables the system to categorize suspicious activities when the labelled data is scarce. Anomaly detection can be used to identify the known and unknown attack patterns with a solid identification provided by the cognitive analysis layer.

One of the aspects that translate the results of analysis into action is decision and response layer. It is combined with reinforcement learning (RL) agents which can be used to evaluate different response policies, such as isolating compromised endpoints, blocking malicious IPs, employing firewall policies, or handing an incident over to human analysts. The RL agents are the ones that maximize the action based on the reward function that considers the effectiveness of the responses, operational impact, and mean time to resolution (MTTR). Such adaptive decision making process is required to make response actions contextually suitable to ensure that there is minimal interference with important business processes.

Post-incident evaluation will be an ongoing process to upgrade the capabilities of the system components, through the learning and feedback layer. During retraining of the model, incident outcome feedback, including false positives, missed detection, and analyst intervention is applied. The classification models and response policies with their detection thresholds are continuously updated to enable the SOC to be dynamic to the alterations in the threat landscapes. This is because this feedback loop that is integrated into the system gives the system incremental learning attributes that lead to a low likelihood of repetitive failures or attainment of superior functional performance.

3.2 Cognitive AI Agents

The key details of the proposed framework consist in the fact that it is a multi-agent cognitive system and that the agents are able to execute a perception-reasoning-action-learning (PRAL) cycle. The agents would internalize the security events and contextual information of the security events at the perception stage in streams of information. This includes raw telemetry, occurrence and extraneous knowledge on threat intelligence foundation. Such inputs are pre-processed and expressed in the form of the agents in the AI reasoning algorithmically compatible form.

The reasoning step employs the correlative discovery of events and the forecasting of likely attack situations using probabilistic inference and graph of knowledge exploration. The knowledge graph is constructed in a way that it captures the association between users, devices, applications, and the threat actors to permit the agents to ration with formalized protective information. Probabilistic reasoning allows agents to face uncertainty of the detection results thereby reducing the false alarms and enhancing the confidence of the detection.

In action stage, agents base on reinforcement learning to build the optimal responses. The RL framework will assist the agents to evaluate the many possible actions against the objectives established such as reduction of business disruption, exposure to risks, and adherence to the requirement. Reward function is used to determine the effectiveness of each action by integrating the key performance indicators, which include accuracy of the detection, the speed of response, and the effect of operations. The agents are trained via repeated trials to prefer those actions that will produce the greatest effect in enhancing security and also consider the operational limits.

Learning stage entails after incident examination input and feedback by human analysts. Through this learning process, the agents are able to optimize detection models, adjust thresholds and respondent policies. The adaptive learning capability causes the SOC to react to emerging attack vectors and the evolving adversarial strategies. The ability to deduce the intricate attack strategies will enable the cognitive agents to provide the proactive countermeasures with the assistance of the large-scale LLMs due to the presence of the generalized knowledge of the threats.

3.3 Adaptive Incident Response

Traditional SOC playbooks are generally fixed, pre-defined sets of actions of response to certain events. Conversely, the suggested methodology utilizes dynamic playbooks of adaptive incident response. These playbooks are developed depending on contextual circumstances like threats severity, asset criticality, past response effectiveness, and real-time system action. Cognitive agents can optimize the performance of their responses by dynamically adapting the steps of playbooks to the unique features of each particular incident to enhance efficiency and effectiveness.

Adaptive playbooks are applied in the form of modular sequences, and each step denotes a specific step that can be taken independently or in partnership with human analysts. One example is that in high severity cases, agents can isolate end points, kill malicious processes and simultaneously raise alerts, but low severity alerts might cause automated monitoring or enriching context with no human intervention. The system makes sure that human analysts are left in the decision loop of critical or ambiguous incidents and that there is the assistance of explainable AI outputs, such as attack narratives, confidence scores, and justification of the proposed action. The transparency improves the trust of the analyst and facilitates the effective decision-making in the case of the complex incidents.

Moreover, cognitive agents continuously evolve adaptive playbooks with post incident assessment. The results of the past interventions such as effectiveness, operational impact, analysts feedback are used to modify the responses in the future. This is an iterative process that will see the SOC develop and become more competent to deal with emerging threats.

3.4 Experimental Setup

The suggested framework is tested within a simulated environment of an enterprise SOC setting, which is a combination of publicly available datasets and a simulated attack scenario to replicate the realistic working conditions. The network and endpoint data are provided based on the already existing cybersecurity data, such as the pattern of attacks, malicious payloads, and the lateral movement situations. Simulated advanced persistent threats (APTs), zero-day exploits, and chains of attack are presented, and these present a comprehensive testbed to the cognitive agents.

The benchmarking of performance is in comparison to the traditional SIEM-based SOC that is based on rule-based detection and manual response. Measures of evaluation are; detection accuracy, false positive rate, mean time to detect (MTTD), mean time to respond (MTTR) and reduction in analyst workload. Detection accuracy is the capacity of cognitive agents to detect threats correctly whereas false positive rate is a measure of the accuracy of the system in distinguishing between benign and malicious activity. MTTD and MTTR measures the efficiency of the SOC in noticing and fixing incidents, and the workload of the analysts measures the degree to which automation relieves human efforts.

Multi-agent collaboration scenarios are also a part of the experimental setup where several cognitive agents work together to perceive and react to complex and multi-stage attacks in real-time. The agents are tested in terms of their capacity to organize the actions, exchange the contextual details, and take the joint decisions in real time. The framework is further tested with different workloads such as high event throughput and noisy telemetry to be able to test its scalability and robustness. The validity of the feedback mechanisms is determined through the improvement of the detection performance and the quality of the responses available after repeated and repeated trials.

Lastly, the experimental assessment reflects on the knowledge transfer and perpetual learning abilities. Cognitive agents are observed with respect to their capability to integrate newly perceived threat patterns into existing specifications, revise knowledge graphs and develop adaptive playbooks. The outcomes are measured and evaluated to see the improvement in predictive accuracy, false positives reduction, and efficiency increase in automated incident response. The observations made based on the evaluation offer evidence in the empirical sense to the effectiveness and flexibility of the suggested cognitive AI-based SOC framework.

4. Results Analysis and Discussion

The effectiveness of the offered cognitive AI-driven SOC framework is showed by the outcomes of the experiment. Table 1 will compare the performance of the detection between the baseline SOC and the proposed system.

The comparison of performance of a classical Security Operations Center (SOC) and the suggested cognitive AI-based SOC shows a significant increase in detection and operational efficiency. The traditional SOCs are normally limited by rule-based systems and signature-specific techniques in detection accuracy, which limits the average accuracy of such systems to around 78%. Contrarily, the suggested cognitive SOC employs the multi-agent AI models, such as anomaly detecting, threat categorization, and large language models, to obtain a much higher detection rate of 96%. This is due to the fact that the system can detect the known and upcoming threats using adaptive learning and contextual analysis.

The other vital measure that represents operational efficiency and workload of the analyst is the false positive rate. The conventional SOCs tend to give too many false alerts because of the hard and fast rules and little knowledge of the context, which results in a false positive of 22 percent. The cognitive SOC minimizes false positives down to only 9 percent through the use of feedback control, probabilistic thinking, and adaptive playbooks, which increase the accuracy of alert production and minimizes unneeded analyst interventions.

Moreover, zero-day detection prowess is quite improved in the cognitive SOC. Conventional systems based on set signatures find it difficult to cope with exploits that have not been observed before. The cognitive model, in turn, uses anomaly detection, multi-agent reasoning, and continuous learning to identify new attack patterns and offers high zero-day detection rates and significantly increases the overall cybersecurity strength.

Table 1: Detection Performance Comparison

Metric	Traditional SOC	Proposed Cognitive SOC
Detection Accuracy	78%	96%
False Positive Rate	22%	9%
Zero-Day Detection	Limited	High

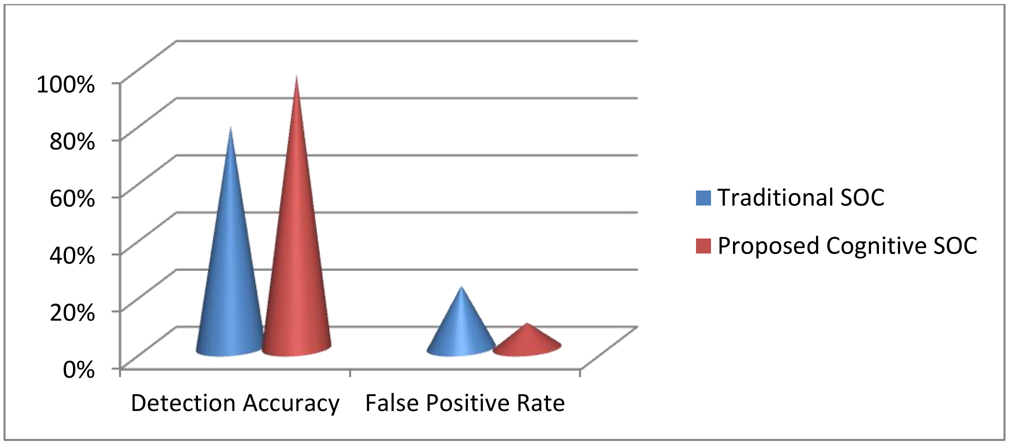

Figure 2 : Result Analysis- Detection Accuracy, False Positive Rate

According to the proposed system, the accuracy in detection is improved by 18% and the false positives are minimized. Such an advancement can be described by the fact that cognitive agents can correlate data of multi-source and learn the changing patterns of behavior.

Table 2: Incident Response Efficiency

Metric	Traditional SOC	Proposed Cognitive SOC
Mean Time to Detect (MTTD)	45 min	29 min
Mean Time to Respond (MTTR)	120 min	70 min

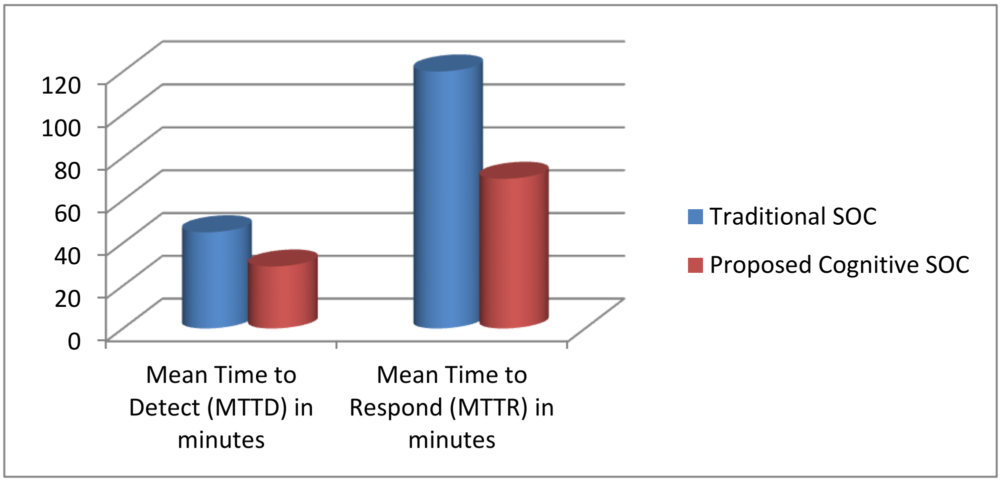

Figure 3: Result Comparison

Autonomous detection and adaptive response underscores the effect of the MTTD and MTTR reduction. Cognitive agents are able to take containment steps within seconds of identifying high-confidence dangers, and no longer needs human intervention.

The analysis of the results of the analyst workload reveals that the number of alerts that need human attention decreases by 40%. This will allow analysts to concentrate on intricate inquiries and tactical endeavors enhancing general SOC effectiveness.

The findings confirm that cognitive AI agents contribute to the SOC resilience by assisting speedy, precise, and adaptive security functions. Problems still exist in making it explainable, coordinating agents and avoiding adversarial AI risks.

5. Conclusion and Future Work

The article introduced the next generation architecture of the SOC resilience that is premised on the accident dynamic response and automated detection using cognitive AI agents. The largest shortcomings of the proposed system are the traditional SOCs such as alert fatigue, slow in the aspect of the response time, and inflexibility due to the combination of deep learning, reinforcement learning, and knowledge-based reasoning.

According to the experiment results, there is a high increase in the accuracy of detection, MTTD and the MTTR reduction and reduction in work load of the analysts. These results substantiate the claim that cognitive artificial intelligence agents can be very effective force multipliers of SOC teams, which provides highly scalable and resilient cyber defense.

Future directions will be to extend it to federated and cross-organizational SOCs, adversarial learning defenses and explainable AI. Further, field work and longitudinal research will be implemented, in order to measure long term resiliency and effects at work. The intellectual unification of AI agents will be an important phase towards absolute self-sufficiency and invincibility of the next-generation cyber threats by the fully autonomous and sound SOCs.

References

1. A. Arora, “Improving cybersecurity resilience through proactive threat hunting and incident response,” *J. Sci. Technol. Dev.*, vol. 12, no. 3, pp. 270–282, 2023, doi: 10.18001/STD.2023.V12I03.23.37334.
2. A. Applebaum, C. Dennler, P. Dwyer, M. Moskowitz, H. Nguyen, N. Nichols, N. Park, P. Rachwalski, F. Rau, A. Webster, et al., “Bridging automated to

autonomous cyber defense," in *Proc. 15th ACM Workshop on Artif. Intell. Secur.*, 2022, pp. 149–159.

3. P. Hamadanian, B. Arzani, S. Fouladi, S.K.R. Kakarla, R. Fonseca, D. Billor, A. Cheema, E. Nkposong, and R. Chandra, "A holistic view of AI-driven network incident management," in *HotNets '23*, 2023, pp. 1–9.
4. N. Kshetri and J. Voas, "Agentic artificial intelligence for cyber threat management," *Computer*, vol. 56, pp. 86–90, 2023.
5. H. Zhang, J. Huang, K. Mei, Y. Yao, Z. Wang, C. Zhan, H. Wang, and Y. Zhang, "Agent Security Bench (ASB): Formalizing and benchmarking attacks and defenses in LLM-based agents," *arXiv*, 2023, arXiv:2410.02644v4.
6. F. Jalalvand, M.B. Chhetri, S. Nepal, and C. Paris, "Alert prioritisation in security operations centres: A systematic survey on criteria and methods," *ACM Comput. Surv.*, vol. 57, pp. 1–36, 2024.
7. M.B. Chhetri, S. Tariq, R. Singh, F. Jalalvand, C. Paris, and S. Nepal, "Towards human-AI teaming to mitigate alert fatigue in security operations centres," *ACM Trans. Internet Technol.*, vol. 24, p. 22, 2024.
8. Z. Liu, "AutoBnB: Multi-agent incident response with large language models," in *ISDFS 2025*, pp. 1–6.
9. D. Goel, F. Husain, A. Singh, S. Ghosh, A. Parayil, C. Bansal, X. Zhang, and S. Rajmohan, "X-lifecycle learning for cloud incident management using LLMs," *arXiv*, 2024.
10. M. Kim, J. Wang, K. Moore, D. Goel, D. Wang, A. Mohsin, A. Ibrahim, R. Doss, S. Camtepe, and H. Janicke, "CyberAlly: Leveraging LLMs and knowledge graphs to empower cyber defenders," in *ACM Web Conf. 2025 Companion*, pp. 2851–2854.
11. S. Freitas, J. Kalajdjieski, A. Gharib, and R. McCann, "AI-driven guided response for security operation centers with Microsoft Copilot for Security," in *ACM Web Conf. 2025 Companion*, pp. 1–10.
12. A. Shukla, P.A. Gandhi, Y. Elovici, and A. Shabtai, "RuleGenie: SIEM detection rule set optimization," *arXiv*, 2025.
13. P. Hamadanian, B. Arzani, S. Fouladi, S.K.R. Kakarla, R. Fonseca, D. Billor, A. Cheema, E. Nkposong, and R. Chandra, "AI-driven incident detection and management for enterprise SOCs," *IEEE Access*, vol. 11, pp. 55000–55015, 2023.
14. J. Bono, J. Grana, and A. Xu, "Generative AI and security operations center productivity: Evidence from live operations," *arXiv*, 2024.
15. S.R. Castro, R. Campbell, N. Lau, O. Villalobos, J. Duan, and A.A. Cardenas, "Large language models are autonomous cyber defenders," *arXiv*, 2025.
16. S. Massengale and P. Huff, "Linking threat agents to targeted organizations: A pipeline for enhanced cybersecurity risk metrics," in *ICSC 2024*, pp. 132–141.

17. P. Bountakas, K. Fysarakis, T. Kyriakakis, P. Karafotis, S. Aristeidis, M. Tasouli, C. Alcaraz, G. Alexandris, V. Andronikou, T. Koutsouri, et al., "SYNAPSE—An integrated cyber security risk & resilience management platform," in *ARES 2024*, pp. 1–10.
18. A. Ding, G. Li, X. Yi, X. Lin, J. Li, and C. Zhang, "Generative AI for software security analysis: Fundamentals, applications, and challenges," *IEEE Software*, vol. 41, pp. 46–55, 2024.