12th INTERNATIONAL SCIENTIFIC CONFERENCE ON DEFENSIVE TECHNOLOGIES
OTEH 2026
Serbia, Tara, 21 - 24 October 2026

COMPARATIVE EVALUATION OF SIEM PLATFORMS IN DETECTING MITRE ATT&CK TECHNIQUES BASED ON FALSE POSITIVE ANALYSIS

MARKO PREMOVIĆ
Serbian Armed Forces, Belgrade, marko.premovic@vs.rs

MLADEN VEINOVIĆ
Singidunum University, Belgrade, mveinovic@singidunum.ac.rs

BOBAN STOJANOVIĆ
Serbian Armed Forces, Belgrade, boban.stojanovic@vs.rs

IGOR PETROVIĆ
Ministry of Defence, Belgrade, igor.petrovic@mod.gov.rs

Abstract: *Security Information and Event Management (SIEM) platforms are essential for centralized security monitoring in Security Operations Centers (SOCs), yet they continue to face a persistent challenge: high volumes of false positive alerts that impair analyst efficiency and delay incident response. This paper presents a comparative evaluation of three SIEM platforms—Wazuh, IBM QRadar, and Splunk Enterprise Security—in detecting selected MITRE ATT&CK techniques. Several attack scenarios, including brute force authentication, network reconnaissance, and malicious command execution, were executed to assess detection accuracy, latency, false positive rates, and alert precision. Results reveal notable differences between open-source and commercial solutions. Enterprise platforms demonstrated superior correlation capabilities and lower false positive rates, while open-source Wazuh offered competitive detection performance at significantly reduced deployment costs. These findings provide practical insights into SIEM performance when detecting adversary behaviors mapped to the MITRE ATT&CK framework, offering guidance for organizations balancing cost, accuracy, and operational efficiency in their security stack selection.*

Keywords: *SIEM, MITRE ATT&CK, False Positives, Threat Detection, Security Monitoring*

1. INTRODUCTION

The increasing complexity of modern information systems and the continuous growth of cyber threats have significantly increased the importance of centralized security monitoring solutions. Organizations generate large volumes of security-related logs originating from endpoints, servers, authentication services, and network devices, making manual analysis of security events increasingly difficult. In such environments, Security Information and Event Management (SIEM) systems have become essential components of Security Operations Centers (SOC), enabling centralized collection, correlation, and analysis of security events in real-time [1].

Despite the widespread adoption of SIEM technologies, one of the major operational challenges remains the high number of false positive alerts generated during event correlation and threat detection processes. Excessive false positives may lead to alert fatigue, increased workload for SOC analysts, and delayed response to legitimate security incidents. Consequently, evaluating the balance between detection capability and false positive generation has become an important aspect of SIEM performance assessment.

In recent years, the MITRE ATT&CK framework has emerged as a standardized knowledge base for modeling adversarial behavior and evaluating defensive mechanisms. ATT&CK-based testing allows researchers and security analysts to simulate realistic attack techniques and objectively measure the detection capabilities of security platforms [2].

Recent SOC studies additionally emphasize that excessive alert volume may contribute to analyst fatigue, operational inefficiency, and delayed incident prioritization. Consequently, modern SIEM evaluation increasingly includes not only detection capability, but also alert quality and operational usability within SOC environments [3].

This paper presents a comparative evaluation of the Wazuh, IBM QRadar, and Splunk Enterprise Security SIEM platforms through the analysis of selected MITRE ATT&CK techniques and false positive events. The research is conducted in a controlled laboratory environment using simulated attack scenarios mapped to ATT&CK techniques, including brute force attacks, network reconnaissance, and suspicious command execution.

The primary contributions of this research are:

1) a controlled ATT&CK-based comparative evaluation of three widely used SIEM platforms under identical attack conditions;

Corresponding author: M. Premovic;
E-mail: marko.premovic@vs.rs;
DOI: 10.5937/OTEH26100P

2) quantitative assessment of false positive generation and detection effectiveness using standardized evaluation metrics;

3) operational analysis of alert quality through Precision-based evaluation of actionable alerts within SOC environments.

The paper is organized as follows. After the introduction, the second chapter presents the theoretical background on SIEM systems, the MITRE ATT&CK framework, and false positive alerts, while the third chapter describes the experimental environment, attack scenarios, and evaluation metrics. The fourth chapter presents the detection results and false positive analysis, followed by the conclusion and directions for future work.

2. THEORETICAL BACKGROUND

2.1 Security Information and Event Management Systems

Security Information and Event Management (SIEM) systems are designed to provide centralized collection, normalization, correlation, and analysis of security events generated within enterprise infrastructures.

Modern SIEM platforms utilize rule-based analytics, behavioral analysis, and threat intelligence integration to identify suspicious activities and generate alerts for SOC analysts [4].

Modern enterprise SIEM platforms additionally support integration with external threat intelligence feeds, enabling improved contextual analysis and prioritization of high-risk events during correlation processes.

However, improperly configured detection rules may generate excessive false positive events, reducing the operational efficiency of security monitoring processes.

Recent comprehensive surveys of SIEM adoption trends confirm that organizations continue to prioritize false positive reduction as a key operational challenge, with correlation rule tuning remaining one of the most resource-intensive SOC activities [13].

2.2 MITRE ATT&CK Framework

The MITRE ATT&CK framework represents a structured knowledge base used for the classification of adversarial tactics and techniques observed during cyber attacks.

One of the primary advantages of the ATT&CK framework is its applicability in detection engineering and security evaluation. By mapping attack scenarios to ATT&CK techniques, security analysts can evaluate the effectiveness of defensive mechanisms and identify potential detection gaps within monitored environments.

The ATT&CK techniques used in this research are presented in Table 1.

Table 1. Selected MITRE ATT&CK Techniques Used in the Experimental Evaluation

Technique	ATT&CK ID
Brute Force	T1110
Network Service Scanning	T1046
Command and Scripting Interpreter	T1059

2.3 False Positive Events

False positive alerts represent security events incorrectly classified as malicious activities by detection systems. High false positive rates may negatively impact SOC operations by increasing analyst workload and reducing the visibility of legitimate threats [5]. Consequently, reducing unnecessary alerts while maintaining high detection accuracy is considered one of the major challenges in SIEM optimization and detection engineering. In practical SOC environments, false positive reduction is commonly achieved through correlation rule tuning, contextual enrichment, and event prioritization mechanisms. Proper rule optimization may significantly improve analyst efficiency without reducing detection coverage [6].

3. EXPERIMENTAL ENVIRONMENT

3.1 Laboratory Setup

The laboratory environment was designed to emulate a small enterprise network consisting of approximately 50 endpoint systems. The infrastructure included Windows workstations, Linux servers, centralized authentication services, and network devices generating routine operational logs. Endpoint activity included user authentication events, administrative operations, file access events, process execution records, and network communications in order to simulate realistic enterprise workloads.

A dedicated Kali Linux system was used for attack execution, while Windows and Ubuntu hosts acted as monitored targets. Security events were collected through native operating system logging mechanisms and forwarded to the evaluated SIEM platforms. Log sources included authentication logs, process creation events, command execution records, and network connection events.

The laboratory topology was intentionally configured to generate both legitimate administrative activity and malicious behavior. This approach enabled observation of SIEM detection performance under conditions that more closely resemble operational SOC environments, where benign and malicious events coexist within the same monitoring infrastructure. Laboratory parameters are shown in Table 2.

Table 2. Laboratory Configurations Parameters

Parameter	Value
Simulated Endpoints	50
Windows Systems	Multiple
Linux Systems	Multiple
SIEM Platforms	Wazuh 4.8, QRadar 7.5, Splunk ES 7.3.2
Attack Workstation	Kali Linux
Log Sources	Auth, Process, Network, System
ATT&CK Techniques Tested	3
Repetition per Scenario	5

To ensure result comparability, all SIEM platforms received identical log streams generated from the same attack scenarios and background network activity. Platform-specific detection rules were configured according to vendor-recommended practices while maintaining comparable monitoring coverage across all evaluated solutions.

3.2 Attack Scenarios

Several attack scenarios mapped to selected MITRE ATT&CK techniques were executed during the experimental phase of this research. The scenarios were designed to represent common adversarial activities frequently encountered in enterprise environments and SOC operations. Each scenario was executed five times under identical laboratory conditions in order to evaluate detection consistency, alert generation behavior, and correlation effectiveness across the analyzed SIEM platforms.

The first scenario simulated credential access attempts using the Hydra tool and was mapped to ATT&CK technique T1110 (Brute Force). Multiple SSH authentication attempts were generated against Linux systems within the laboratory environment using valid and invalid credential combinations. The objective of this scenario was to evaluate the ability of SIEM platforms to detect repeated authentication failures and identify potential brute-force activity while minimizing unnecessary alerts generated by legitimate authentication events.

The second scenario focused on network reconnaissance activities corresponding to ATT&CK technique T1046 (Network Service Scanning). Reconnaissance was performed using the Nmap tool [8] against multiple hosts within the monitored network segment. Various scanning techniques were executed in order to generate network discovery events and service enumeration activity. This scenario was selected because reconnaissance activities often produce event patterns that resemble legitimate administrative operations, making accurate detection particularly challenging for SIEM platforms.

The third scenario simulated malicious command execution behavior associated with ATT&CK technique T1059 (Command and Scripting Interpreter). PowerShell commands commonly observed in attack investigations were executed on Windows systems in order to generate process creation events, command execution logs, and related security telemetry. The objective was to evaluate the capability of each SIEM solution to identify suspicious command execution activity while maintaining acceptable alert precision.

To provide realistic operating conditions, attack execution was performed while legitimate background activity was simultaneously generated throughout the laboratory environment. Authentication events, file access operations, system administration tasks, and normal network communications were continuously produced by the simulated endpoints. This approach enabled evaluation of SIEM performance under conditions that more closely resemble operational enterprise networks rather than isolated attack demonstrations.

Each attack scenario was repeated five times for every evaluated SIEM platform. During every execution, detection events, generated alerts, false positive occurrences, and detection latency were recorded. The collected observations were subsequently aggregated and used for the calculation of the evaluation metrics presented in this study.

3.3 Evaluation Metrics

The performance evaluation of the three SIEM platforms (Wazuh, IBM QRadar, Splunk Enterprise Security) was based on quantitative metrics derived from controlled attack simulations mapped to MITRE ATT&CK techniques T1110 (Brute Force), T1046 (Network Service Scanning), and T1059 (Command and Scripting Interpreter). Each attack scenario was repeated five times per platform, with simultaneous generation of legitimate background activity across 50 simulated endpoints to emulate realistic SOC conditions.

3.3.1 Determination of Event Classification Outcomes

For each executed attack instance and each background activity event, security events generated by the SIEM platforms were classified into four mutually exclusive categories:

True Positive (TP) -- An alert was generated, and the event corresponded to an actual executed attack technique (e.g., a Hydra brute-force SSH attempt correctly detected as T1110).

False Positive (FP) -- An alert was generated, but the event was not part of any attack; it originated exclusively from legitimate administrative or background activity (e.g., a scheduled PowerShell script flagged as suspicious command execution).

True Negative (TN) -- No alert was generated, and the event was indeed benign (e.g., a normal user login or routine file access).

False Negative (FN) -- No alert was generated, although an attack was successfully executed (e.g., an Nmap network scan that remained undetected).

Classification was performed by correlating SIEM alert logs with ground truth attack logs collected from the Kali Linux attack workstation and from target endpoints (Windows and Linux system logs). Each event was manually verified by the research team to eliminate ambiguity.

The total event set for each SIEM platform consisted of approximately 1000 benign background events (authentication, process creation, file access, network connections) and 120 attack related events (5 repetitions × 3 techniques × approximately 8 individual attack actions per technique). Thus, the total number of evaluated events per platform was 1120, distributed as TP + FP + TN + FN = 1120.

3.3.2 Detection Time Measurement

Detection Time (detection latency) was defined as the interval between the moment an attack action was completed on the target system and the moment the SIEM platform generated a corresponding alert in its console.

Detection latency was measured as the elapsed time between attack execution and alert generation within the corresponding SIEM platform. Measurements were collected for all true positive detections and averaged across all attack scenarios. The final "Avg Detection Time" reported in Table 4 represents the arithmetic mean value per platform.

3.3.3 Calculation of Accuracy and False Positive Rate (FPR)

Detection Accuracy was calculated using the standard formula:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (1)$$

Accuracy represents the proportion of all evaluated events (both malicious and benign) that were correctly classified by the SIEM.

False Positive Rate (FPR) was defined as:

$$FPR = \frac{FP}{FP + TN} \quad (2)$$

FPR measures the ratio of benign events that were incorrectly flagged as malicious relative to all actual benign events. A lower FPR indicates better alert precision and reduced analyst fatigue.

Alert Relevance Ratio (ARR) is proposed as a complementary SOC-oriented metric:

$$ARR = \frac{TP}{TP + FP} \quad (3)$$

Precision represents the proportion of generated alerts that correspond to actual malicious activity. Higher Precision values indicate a greater proportion of actionable alerts requiring analyst investigation.

From an operational SOC perspective, Precision is additionally discussed in this paper as the Alert Relevance Ratio (ARR), emphasizing the practical value of alert quality and analyst workload reduction.

3.3.4 Raw Detection Statistics

Based on the classification methodology described in Section 3.3.1, the raw counts of TP, FP, TN, and FN were aggregated for each SIEM platform across all attack repetitions and background events. Table 3 presents these raw values.

Table 3. Raw Detection Statistics

Platform	TP	FP	TN	FN	Total Events
Wazuh	84	196	804	36	1120
QRadar	90	88	912	30	1120
Splunk	93	67	933	27	1120

Approximately 1000 benign events were generated through routine administrative activity, authentication operations, process execution, and normal network communications across the laboratory environment.

3.3.5 Derived Performance Metrics

Using the raw data from Table 3 and the equations (1) - (3), the following performance metrics were calculated for each SIEM platform.

Avg Detection Time – mean latency in seconds (measured only for TP events).

Table 4. Comparative Performance Metrics Across SIEM Platforms

Platform	Recall	(ARR)	F1	(FPR)	Accuracy	Avg Detection Time (s)
Wazuh	70%	30%	0.42	19.6%	79.3%	4.3 ± 1.2
QRadar	75%	50.6%	0.60	8.8%	89.5%	2.1 ± 0.5
Splunk	77.5%	58.1%	0.66	6.7%	91.6%	1.8 ± 0.4

Values represent the arithmetic mean across five attack repetitions. Standard deviations for detection time were as follows: Wazuh ±1.2 s, QRadar ±0.5 s, Splunk ±0.4 s, indicating consistent detection latency across attack repetitions.

The obtained metrics indicate that enterprise SIEM platforms (QRadar and Splunk) demonstrated significantly higher detection accuracy, substantially lower false positive rates, better alert relevance, and lower detection latency compared to the open-source platform Wazuh. However, Wazuh maintained acceptable detection performance with considerably reduced deployment costs, making it a viable option for organizations with limited resources.

High false positive rates were observed primarily during reconnaissance-related activities (T1046) due to the

similarity between legitimate administrative network scans and malicious scanning behavior. Such conditions may negatively affect SOC workflows by increasing alert fatigue and reducing analyst efficiency. The Alert Relevance Ratio (ARR) further highlights that in the case of Wazuh, only 30% of generated alerts were truly malicious, whereas Splunk achieved nearly 60% relevance.

4. RESULTS AND ANALYSIS

4.1 MITRE ATT&CK Detection Results

Detection performance is reported primarily using recall, precision, and their harmonic mean, since these metrics more faithfully reflect the operational reality of a SOC than accuracy under the present class distribution. The evaluation set contains 120 attack events against approximately 1000 benign events, so accuracy is dominated by the correct classification of benign traffic and overstates platform effectiveness. Splunk Enterprise Security achieves 91.6% accuracy, yet its recall is 77.5%, meaning roughly one in five attack events is missed. Recall follows the same ordering across platforms, at 70.0% for Wazuh, 75.0% for QRadar, and 77.5% for Splunk, indicating that detection coverage is broadly comparable and that the platforms separate far more sharply on alert quality than on raw coverage. Precision, reported here as the Alert Relevance Ratio, exposes that separation directly. Only 30.0% of Wazuh alerts correspond to genuine attacks, against 50.6% for QRadar and 58.1% for Splunk, so the open-source configuration imposes a substantially heavier triage burden for a similar level of coverage. The F1 score summarises both axes as 0.42, 0.60, and 0.66, respectively, confirming that Splunk dominates QRadar, which in turn dominates Wazuh, since each platform achieves both higher recall and higher precision than the platform below it. This ordering reinforces the central observation of this study, that alert relevance rather than detection coverage is the dimension on which these platforms most clearly diverge.

This finding is consistent with the base-rate fallacy problem in intrusion detection, which Axelsson [10] demonstrated nearly two decades ago, yet it remains a persistent challenge in SOC operations.

The performed experiments demonstrated that all analyzed SIEM platforms were capable of detecting the majority of executed ATT&CK techniques. However, differences were observed regarding correlation quality, alert precision, and detection latency.

Similar ATT&CK-based SIEM evaluation approaches have recently been adopted in enterprise threat detection studies [9].

In Table 5, "Partial" detection for Wazuh against T1046 (Network Service Scanning) indicates that while the platform successfully detected certain Nmap scan types (e.g., TCP SYN scan), other scan variants (e.g., UDP scan) remained undetected due to the absence of specific correlation rules in the default Wazuh configuration. In contrast, QRadar and Splunk demonstrated comprehensive detection across all scan types.

Table 5. Detection Results of MITRE ATT&CK Techniques Across Analyzed SIEM Platforms

Technique	Wazuh 4.8	QRadar 7.5	Splunk Enterprise Security 7.3.2
T1110	Detected	Detected	Detected
T1046	Partial	Detected	Detected
T1059	Detected	Detected	Detected

The obtained results indicate that enterprise platforms demonstrated improved event correlation and lower alert redundancy compared to the open-source solution [7].

4.2 False Positive Analysis

False positive analysis showed notable differences between the analyzed SIEM platforms. During network scanning scenarios, the open-source platform generated a larger number of low-priority alerts, while enterprise platforms applied more precise correlation mechanisms that reduced unnecessary event generation.

The Precision-Recall curves in Figure 1 illustrate the trade-off between detection precision and recall across different classification thresholds. Splunk Enterprise Security consistently achieves higher precision than QRadar for the same recall level, while Wazuh exhibits the lowest precision across the entire recall range. This confirms that commercial platforms provide better alert quality, which directly reduces the number of false positives requiring analyst review.

High false positive rates were observed primarily during reconnaissance-related activities due to the similarity between legitimate administrative operations and suspicious scanning behavior. Such conditions may negatively affect SOC workflows by increasing alert fatigue and reducing analyst efficiency.

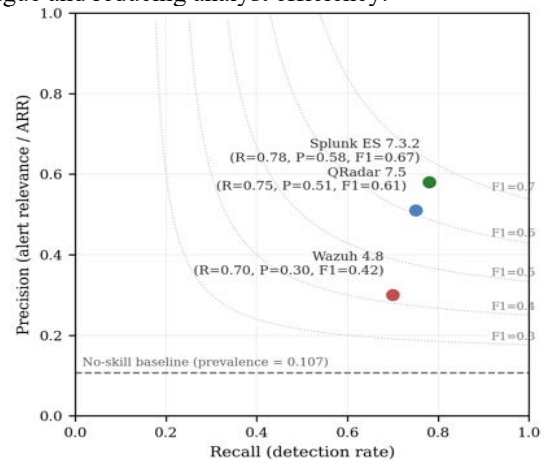


Figure 1. Precision-Recall Curves for the Evaluated SIEM Platforms

The operational impact of false positive rates can be illustrated through a typical SOC scenario. Assuming a daily volume of 1,000 benign events requiring analysis, Wazuh's FPR of 19.6% would generate approximately

196 false positive alerts per day, whereas Splunk's FPR of 6.7% would generate only 67 false positives. This difference of 129 additional alerts per day contributes directly to analyst fatigue and may delay the investigation of true positive incidents. Consequently, the choice of SIEM platform has measurable implications for SOC staffing requirements and mean time to respond (MTTR). The operational impact of false positive alerts on SOC workflows has been extensively documented in recent studies [11].

Beyond the numerical false positive rate, an important observation was the operational impact of alert clustering. Some reconnaissance-related false positives appeared as isolated events, while others triggered multiple correlated alerts for the same activity. Although this did not directly affect detection accuracy, it increased the number of events requiring analyst review. This observation aligns with recent qualitative research by Alahmadi et al. [12], who found that SOC analysts perceive the majority of security alerts as non-actionable, contributing to alert fatigue and operational inefficiency.

4.3 Performance Evaluation

Performance analysis included resource utilization and average detection latency during attack execution. The enterprise platforms demonstrated improved scalability and faster event correlation under higher event volumes, while the open-source platform required lower deployment costs and hardware resources.

The comparative analysis indicates that the selection of a SIEM platform strongly depends on organizational requirements, available infrastructure, and operational priorities.

It should be noted that the presented results were obtained within a controlled laboratory environment with a limited number of attack scenarios and endpoints. Real-world enterprise infrastructures may introduce additional variables, including encrypted traffic, heterogeneous logging formats, and significantly larger event volumes.

Another limitation of the presented evaluation is the relatively small number of ATT&CK techniques included in the testing process. While the selected techniques represent commonly observed adversarial behaviors, modern attack campaigns often combine multiple tactics across different stages of the intrusion lifecycle. Consequently, the obtained results primarily reflect platform performance under focused attack scenarios rather than full attack-chain simulations. Future evaluations may benefit from incorporating multi-stage attack emulation in order to assess SIEM correlation capabilities under more complex operational conditions.

5. CONCLUSION

This paper presents a comparative evaluation of the Wazuh, IBM QRadar, and Splunk Enterprise Security SIEM platforms through the analysis of selected MITRE ATT&CK techniques and false positive events. The conducted experiments demonstrated that all analyzed platforms were capable of detecting common adversarial activities, including brute force attacks, network reconnaissance, and suspicious command execution.

The obtained results showed that enterprise SIEM solutions provided improved event correlation and reduced false positive rates, while the open-source platform offered competitive detection capabilities with significantly lower deployment costs. Furthermore, the research highlighted the operational impact of false positive alerts on SOC workflows [9] and emphasized the importance of correlation rule optimization in modern security monitoring environments [3][6].

ARR is proposed as a complementary indicator for future SIEM benchmarking studies.

The study also indicates that alert relevance and correlation quality may represent equally important evaluation criteria as traditional detection metrics when assessing SIEM effectiveness in operational SOC environments.

Based on the obtained results, organizations with limited budgets and smaller infrastructures may benefit from Wazuh's competitive detection performance at reduced costs, while large enterprises requiring high precision and low latency should consider commercial solutions such as Splunk Enterprise Security or IBM QRadar.

Future research may include larger-scale infrastructures, additional ATT&CK techniques, and integration of machine learning approaches for adaptive SIEM correlation and false positive reduction.

The findings suggest that alert quality metrics may provide greater operational insight than detection accuracy alone. This is consistent with the base-rate fallacy problem identified by Axelsson [10], who demonstrated that accuracy becomes misleading when attacks constitute a small fraction of total events, as is typical in enterprise environments.

References

[1] K. Kent, M. Souppaya, "Guide to Computer Security Log Management," NIST Special Publication 800-92, National Institute of Standards and Technology, 2006.
Available: https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-92.pdf DOI: 10.6028/NIST.SP.800-92

[2] MITRE ATT&CK Team, “MITRE ATT&CK: Design and Philosophy,” MITRE Corporation, 2020. Available: https://attack.mitre.org/

[3] Splunk Threat Research Team, “The State of Security Operations: Alert Fatigue and SOC Efficiency,” Splunk Research Report, 2024. Available: https://www.splunk.com/

[4] S. Peisert, M. Bishop, "Improving Security Event Correlation in Modern SOC Environments," IEEE Security & Privacy, vol. 20, no. 3, pp. 44-52, 2022. DOI: 10.1109/MSEC.2022.3154325

[5] A. Herrero, E. Corchado, "False Positive Reduction in Intrusion Detection Systems," International

Symposium on Distributed Computing and Artificial Intelligence,pp. 253-258, 2009.
DOI: 10.1007/978-3-642-02481-8_34

[6] Elastic Security Labs, “Detection Engineering and SIEM Rule Tuning Practices,” Elastic Security Research, 2023. Available: https://www.elastic.co/security-labs/

[7] G. Lyon, “Nmap Network Scanning: The Official Nmap Project Guide to Network Discovery and Security Scanning,” Insecure.Com LLC, 2009.

[8] A. Ahmed, K. Kim, "Evaluation of SIEM Platforms for Enterprise Threat Detection Using MITRE ATT&CK," Computers & Security, vol. 128, 2023. DOI: 10.1016/j.cose.2023.103185

[9] M. C. D'Amico, J. Whitley, "The Real Cost of Alert Fatigue in Security Operations Centers," IEEE Security & Privacy, 2023. DOI: 10.1109/MSEC.2023.3289472

[10] S. Axelsson, "The Base-Rate Fallacy and the Difficulty of Intrusion Detection," *ACM Transactions on Information and System Security*, vol. 3, no. 3, pp. 186-205, 2000, doi: 10.1145/357830.357849.

[11] B. A. Alahmadi, L. Axon, and I. Martinovic, "99% False Positives: A Qualitative Study of SOC Analysts' Perspectives on Security Alarms," in Proc. 31st USENIX Security Symposium, 2022, pp. 2783-2800.

[12] G. González-Granadillo, S. González-Zarzosa, and R. Díaz, "Security Information and Event Management (SIEM): Analysis, Trends, and Usage in Critical Infrastructures," *Sensors*, vol. 21, no. 14, art. 4759, 2021, doi: 10.3390/s21144759.