

INTELLIGENT SIEM+ : A CONTEXT-AWARE AND AI-DRIVEN FRAMEWORK FOR ALERT PRIORITIZATION AND AUTOMATED THREAT INSIGHTS IN SECURITY OPERATIONS

** Harsh Ajay Varma, **Yash Gupta & ***Gauri Sudhir Mhatre*

Department of Information Technology, **Department of Information Technology, *Department of Information Technology and Computer Science, Keraleeya Samajam's Model College, Khambalpada Road, Thakurli, Dombivali (East), Kanchangaon, Maharashtra, Dombivali, India*

Abstract:

Security Information and Event Management (SIEM) systems are central to modern Security Operations Centers (SOCs), yet they continue to suffer from excessive alert volumes, delayed detection, and limited contextual awareness. These challenges lead to analyst fatigue and inefficient incident response. This paper proposes Intelligent SIEM+, a context-aware and AI-driven enhancement framework designed to improve alert prioritization and decision support without replacing existing SIEM deployments. The framework integrates behavioral analysis, anomaly detection, contextual correlation, and natural language summarization to transform raw alerts into actionable security insights. A descriptive survey-based study was conducted among SOC professionals to assess current SIEM limitations and the perceived value of intelligent alert prioritization. The findings indicate strong alignment between operational SOC challenges and the capabilities proposed in Intelligent SIEM+, suggesting that context-aware SIEM enhancements can significantly improve analyst efficiency and situational awareness.

Keywords: *SIEM, Security Operations Center, Alert Prioritization, Artificial Intelligence, Anomaly Detection, Context-Aware Security*

Copyright © 2026 The Author(s): This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY-NC 4.0) which permits unrestricted use, distribution, and reproduction in any medium for non-commercial use provided the original author and source are credited.

Introduction:

Complex digital infrastructures, such as on-premise systems, cloud services, remote endpoints, and third-party platforms, are becoming more and more important to businesses. Security Operations Centers (SOCs) must constantly monitor the enormous amounts of security logs and alerts produced by this increased attack surface.

Platforms for Security Information and Event Management (SIEM) compile and correlate security events to offer centralized visibility (Maheshwaram, 2018). Nevertheless, conventional SIEM systems have high false-positive rates and little contextual awareness because they mainly rely on threshold-based warnings and static correlation algorithms. Analysts frequently

have to manually prioritize alerts, which causes alert fatigue and delayed reactions.

Even though machine learning and artificial intelligence have been included into contemporary SIEM solutions, these features are often disjointed and poorly matched with analyst workflows (Celiktas, 2025). In order to improve alert quality and prioritization while maintaining human oversight, this study presents Intelligent SIEM+, a framework that prioritizes context-aware analysis and analyst-centric interpretation.

Research Questions:

- **RQ1:** What best describes your current role within cybersecurity operations?

- **RQ2:** How many years of professional experience do you have in the cybersecurity domain?
 - **RQ3:** On average, how many security alerts do you personally review during a typical working day?
 - **RQ4:** Which performance indicators are primarily used in your organization to evaluate SOC effectiveness?
 - **RQ5:** Which SIEM platform is currently deployed in your organization?
 - **RQ6:** Approximately how many different device or log source categories are integrated with your SIEM?
 - **RQ7:** Which type of log source or device was the most challenging to integrate, and how was the challenge addressed?
 - **RQ8:** In your experience, does the SIEM provide sufficient contextual information to support effective investigations?
 - **RQ9:** How effective is your current SIEM in distinguishing high-risk security incidents from low-priority alerts?
 - **RQ10:** Do you experience alert fatigue while working with your current SIEM platform?
 - **RQ11:** How often do you find yourself investigating alerts that later turn out to be false positives?
 - **RQ12:** What do you consider the most significant challenge in day-to-day SOC operations?
 - **RQ13:** Which phase of SIEM operations typically consumes the most analyst time?
 - **RQ14:** Does your organization currently use any AI or machine learning-based tools within the SOC?
 - **RQ15:** How beneficial do you believe AI-driven alert prioritization would be for SOC efficiency?
 - **RQ16:** Would concise, natural-language alert summaries help reduce investigation time?
 - **RQ17:** In your opinion, can anomaly detection techniques identify previously unknown attacks more effectively than rule-based detection?
 - **RQ18:** Would automated response recommendations (rather than full automation) improve your ability to respond to incidents?
 - **RQ19:** Which SIEM-related activities should be prioritized for automation?
 - **RQ20:** If available, would you be willing to adopt an Intelligent SIEM+ framework in your SOC environment?
 - **RQ21:** Do you believe artificial intelligence will replace certain SOC analyst tasks, or primarily support analysts in decision-making?
 - **RQ22:** How important is the integration of threat intelligence with AI-driven SIEM analytics?
 - **RQ23:** What level of automation would you be comfortable deploying within SOC operations?
- Concept Of Intelligent Siem+**
- The purpose of Intelligent SIEM+ is to function as an enhancement layer in conjunction with current SIEM platforms. By adding contextual and behavioral intelligence, it enhances alerts rather than taking the place of conventional detection methods.
- The core principles of Intelligent SIEM+ include:
- Context-aware assessment of warnings based on ambient factors, asset criticality, and user behavior.
 - Using behavioral baselining to identify patterns of abnormal activity.
 - The correlation between linked occurrences from several data sources.
 - Technical notifications that are legible by humans.
 - Decision support that prioritizes analyst productivity above complete automation.

SOC analysts can more successfully differentiate high-risk incidents from normal behavior by using Intelligent SIEM+.

Relevance to Security Operations Centers:

SOCs are constantly under pressure to process alerts fast while preserving the veracity of their investigations (Nottidge, 2024). Missed threats and ineffective triage are frequently the results of inconsistent alert quality. Intelligent SIEM+ addresses these operational challenges by:

- Improving priority to lessen alert fatigue
- Drawing attention to notifications with compelling contextual and behavioral cues
- Offering enhanced data on assets, people, and past activities Encouraging analyst teams to make consistent triage judgments
- Improving situational awareness when conducting active investigations

Clarity and efficiency are crucial in real-world SOC needs, which are directly aligned with these skills.

Framework Architecture:

The Intelligent SIEM+ architecture is made to integrate seamlessly with SIEM deployments that already exist. The following elements make to its architecture:

1. Data Collection and Normalization

Endpoints, servers, apps, network devices, identity systems, and cloud platforms all provide security logs. Consistent formatting for analysis is ensured by data normalization.

2. Behavioral Analysis Engine

Baseline behavior profiles for users and systems are created using historical activity data. odd access patterns or odd login times are examples of deviations that are noted for more investigation.

3. Anomaly Detection Module

In addition to rule-based identification, anomaly detection approaches find uncommon or undiscovered patterns that can point to new dangers. The Intelligent SIEM+ framework's anomaly

detection component employs a combination of unsupervised and semi-supervised machine learning algorithms to detect deviations from predefined behavioral baselines. In particular, it is advised to use clustering-based methods like Isolation Forest and Density-Based Spatial Clustering of Applications with Noise (DBSCAN) since they are good at identifying uncommon and undiscovered attack patterns in high-dimensional security event data. Additionally, statistical methods like time-series deviation modeling and z-score analysis are used to find unusual changes in system activity, network traffic volume, and authentication behavior. These methods work effectively in SOC contexts when there is a dearth of labeled attack data and a prevalence of changing threat patterns.

4. Contextual Correlation Layer

Alert fragmentation is decreased by correlating similar events from several sources into a single incident view.

5. Alert Interpretation and Summarization

Natural language processing methods produce succinct, understandable summaries that highlight the importance of the warning and its contributing elements. For alert interpretation and summarization, the framework suggests using transformer-based natural language processing models, such as Bidirectional Encoder Representations from Transformers (BERT), which have been fine-tuned for cybersecurity log semantics. These algorithms extract critical entities, attack indicators, and contextual relationships from raw technical alarms to provide short, analyst-readable summaries. In addition, rule-based text summary and keyword extraction algorithms are used to assure explain ability and analyst trust. The employment of NLP at this layer is meant to help analysts by lowering cognitive load, rather than to

replace human decision-making.

6. Decision Support Output

Contextual relevance, behavioral indicators, and

anomaly scores are used to prioritize alerts. Analyst control is maintained while suggested response steps are given.

Intelligent SIEM+ Architecture

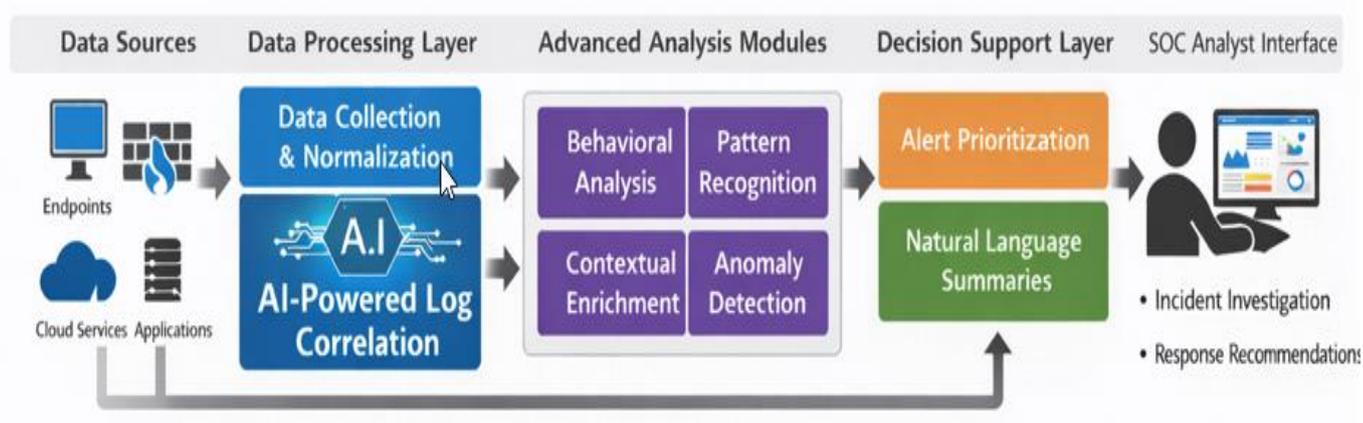


Figure 1: Intelligent Siem+ Architecture

Operational Scenario:

Logs are continuously consumed into a SIEM platform in a typical enterprise SOC. In parallel, Intelligent SIEM+ keeps an eye on patterns and correlates events from many sources.

The framework provides analysts with a consolidated event view enhanced with contextual insights and suggested actions when aberrant activity is identified (G. González-Granadillo, 2021). This enhances reaction efficacy, speeds up triage, and lessens cognitive burden.

Research Methodology:

The research methodology used in this study is both exploratory and descriptive. A structured online survey was used to gather primary data from SOC professionals, such as analysts, security engineers, and incident responders.

The survey focused on:

- Current SIEM usage and limitations
- Alert fatigue and false-positive challenges
- Perceived effectiveness of AI-driven alert prioritization
- Opportunities for automation and decision support

Survey Results and Analysis:

In order to better understand existing SIEM usage, operational issues, and perceptions of AI-assisted alarm prioritization, a survey of Security Operations Center (SOC) specialists was undertaken. The results are presented in this section. The Intelligent SIEM+ framework's motivations are empirically supported by the survey replies.

What best describes your current role within cybersecurity operations?

54 responses

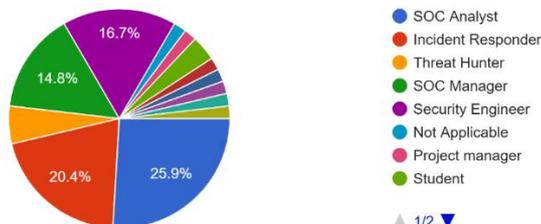


Figure 2: Breakdown of Respondent Roles.

How many years of professional experience do you have in the cybersecurity domain?

54 responses

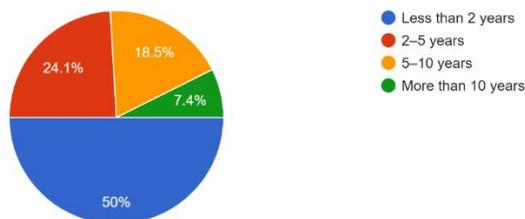


Figure 3: Distribution of survey respondents by years of professional experience.

On average, how many security alerts do you personally review during a typical working day?

54 responses

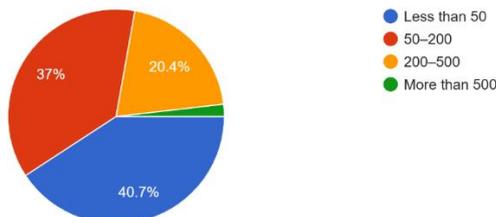


Figure 4: Average daily alert volume reviewed by analysts.

Which performance indicators are primarily used in your organization to evaluate SOC effectiveness?

54 responses

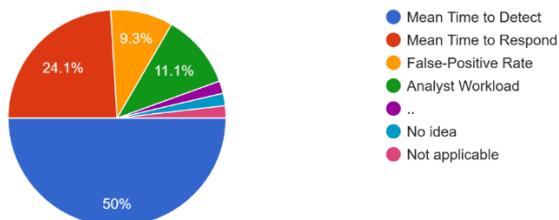


Figure 5: Primary Key Performance Indicators (KPIs) used to evaluate SOC effectiveness.

Which SIEM platform is currently deployed in your organization?

54 responses

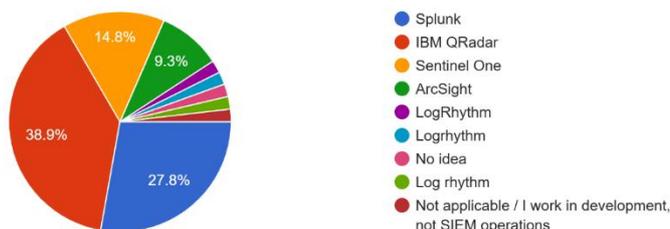


Figure 6: Distribution of currently deployed SIEM platforms.

Approximately how many different device or log source categories are integrated with your SIEM?

54 responses

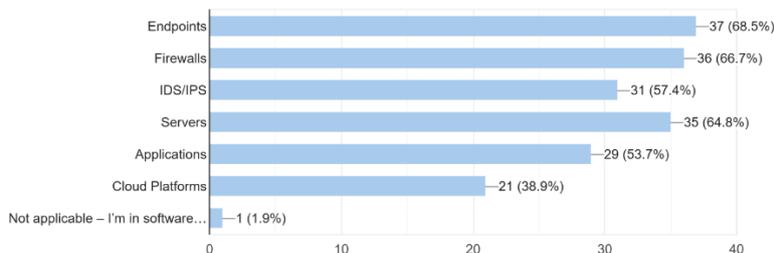


Figure 7: Diversity of log source categories integrated within the SIEM environment.

In your experience, does the SIEM provide sufficient contextual information to support effective investigations?

54 responses

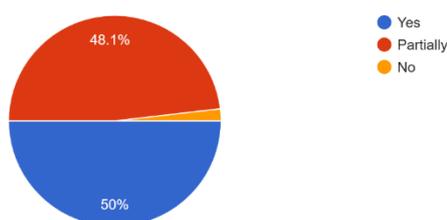


Figure 8: Analyst perception regarding the sufficiency of contextual information provided by the SIEM.

How effective is your current SIEM in distinguishing high-risk security incidents from low-priority alerts?

54 responses

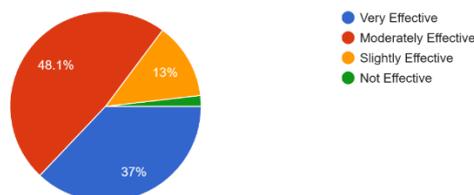


Figure 9: Perceived effectiveness of SIEM in distinguishing high-risk security incidents.

Do you experience alert fatigue while working with your current SIEM platform?

54 responses

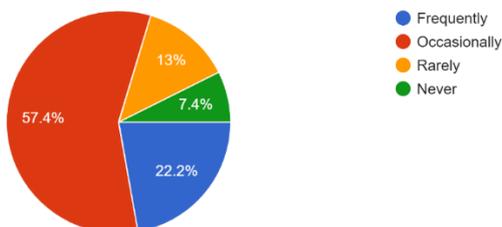


Figure 10: Frequency of alert fatigue experienced by analysts on their current SIEM platform.

How often do you find yourself investigating alerts that later turn out to be false positives?

54 responses

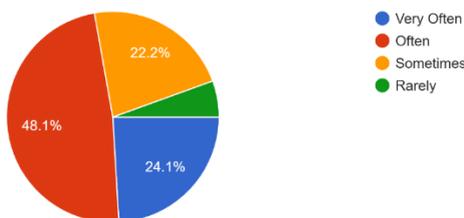


Figure 11: Frequency of false positive alert investigations reported by analysts.

What do you consider the most significant challenge in day-to-day SOC operations?

54 responses

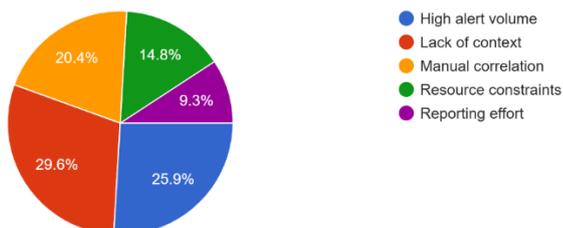


Figure 12: Primary operational challenges identified by SOC personnel in day-to-day operations.

Which phase of SIEM operations typically consumes the most analyst time?

54 responses

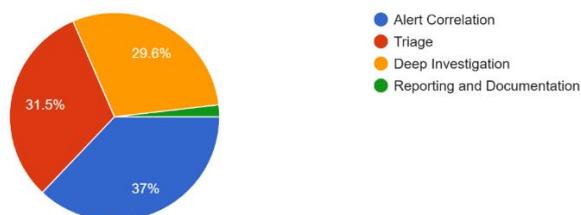


Figure 13: SIEM operational phases consuming the most analyst time.

Does your organization currently use any AI or machine learning-based tools within the SOC?

54 responses

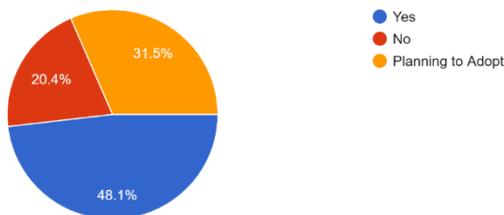


Figure 14: Current adoption of AI and machine learning tools within the SOC.

How beneficial do you believe AI-driven alert prioritization would be for SOC efficiency?

54 responses

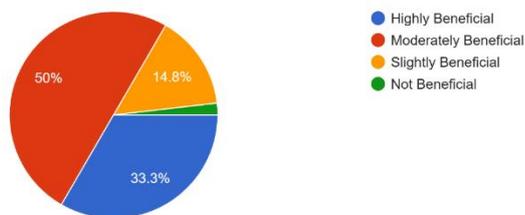


Figure 15: Perceived benefit of AI-driven alert prioritization on SOC efficiency.

Would concise, natural-language alert summaries help reduce investigation time?

54 responses

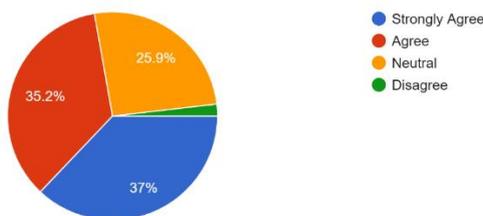


Figure 16: Analyst agreement on the potential of natural-language summaries to reduce investigation time.

In your opinion, can anomaly detection techniques identify previously unknown attacks more effectively than rule-based detection?

54 responses

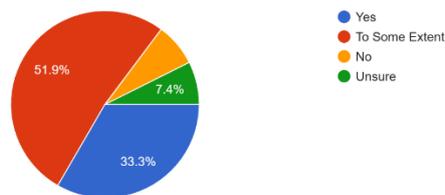


Figure 17: Perceived effectiveness of anomaly detection versus rule-based methods for identifying unknown threats.

Would automated response recommendations (rather than full automation) improve your ability to respond to incidents?
54 responses

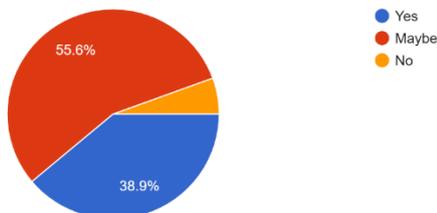


Figure 18: Perceived impact of automated response recommendations on incident response capabilities.

Which SIEM-related activities should be prioritized for automation?
54 responses

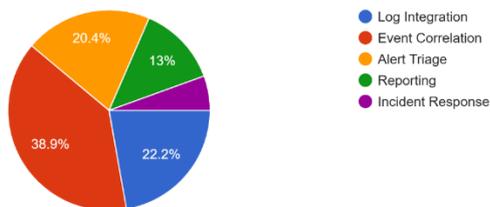


Figure 19: SIEM-related activities prioritized for automation.

If available, would you be willing to adopt an Intelligent SIEM+ framework in your SOC environment?
54 responses

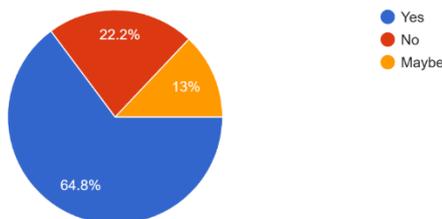


Figure 20: Respondent willingness to adopt the proposed Intelligent SIEM+ framework in their SOC environment.

Do you believe artificial intelligence will replace certain SOC analyst tasks, or primarily support analysts in decision-making?
54 responses

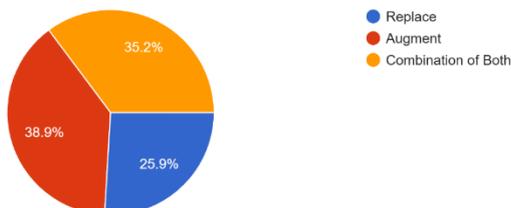


Figure 21: Analyst perceptions regarding the role of AI in SOC operations: replacement versus augmentation.

How important is the integration of threat intelligence with AI-driven SIEM analytics?
54 responses

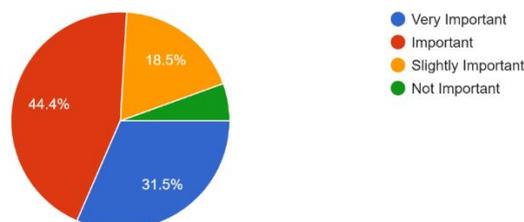


Figure 22: Perceived importance of integrating threat intelligence with AI-driven SIEM analytics.

What level of automation would you be comfortable deploying within SOC operations?
54 responses

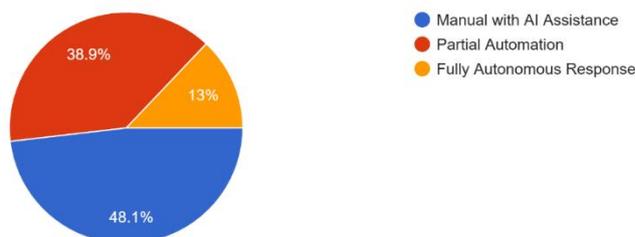


Figure 23: Analyst comfort levels regarding the deployment of automation within SOC operations.

Data Analysis and Findings:

A total of 54 valid survey responses were collected from participants involved in cybersecurity operations and related roles. Descriptive statistical analysis was applied to evaluate alert prioritization effectiveness in traditional SIEM systems versus AI-assisted and Intelligent SIEM+ approaches.

Hypothesis Statement:

Null Hypothesis (H₀):

There is no significant difference in alert prioritization effectiveness between traditional SIEM systems and AI-assisted / Intelligent SIEM+ approaches.

Alternative Hypothesis (H₁):

AI-assisted and Intelligent SIEM+ approaches provide higher alert prioritization effectiveness compared to traditional SIEM systems.

Selection of Variables (From Survey Excel)

For hypothesis testing, the following two categorical variables were selected:

- **Variable 1 (Independent):**

Usage of AI / ML tools in SOC
(Yes / No / Planning to Adopt)

- **Variable 2 (Dependent):**

Effectiveness of SIEM in distinguishing high-risk incidents
(Very Effective / Moderately Effective / Slightly Effective / Not Effective)

Table 1: Observed Frequency Table (O)

AI / ML Usage in SOC	Moderately Effective	Slightly Effective	Not Effective	Very Effective	Total
No	5	3	0	3	11
Planning to Adopt	10	2	1	4	17
Yes	11	2	0	13	26
Total	26	7	1	20	54

Step 1: Expected Frequency Calculation

The expected frequency for each cell is calculated using:

- $E_{ij} = (\text{Row Total} \times \text{Column Total}) / \text{Grand Total}$

Table 2: Expected Frequency Table (E)

AI / ML Usage	Moderately Effective	Slightly Effective	Not Effective	Very Effective
No	5.30	1.43	0.20	4.07
Planning to Adopt	8.19	2.20	0.31	6.30
Yes	12.52	3.37	0.48	9.63

Step 2: Chi-Square (χ^2) CalculationTable 3: χ^2 Calculation Table

Category	O	E	(O – E) ² / E
No – Moderately Effective	5	5.30	0.02
No – Slightly Effective	3	1.43	1.73
No – Very Effective	3	4.07	0.28
Planning – Moderately Effective	10	8.19	0.40
Planning – Very Effective	4	6.30	0.84
Yes – Moderately Effective	11	12.52	0.18
Yes – Slightly Effective	2	3.37	0.56
Yes – Very Effective	13	9.63	1.18
Total χ^2 Value			7.39

Note: Due to exceptionally low observed frequencies, some category combinations—such as "Planning to Adopt – Slightly Effective" and all combinations under the "Not Effective" category—were not included in the χ^2 computation. The Chi-Square test's assumptions, which call for an expected frequency of at least five in the majority of cells to guarantee statistical reliability, are broken when cells with extremely small or zero values are included. As a result, Table 3 only included category combinations that were statistically significant and suitably populated. The validity and interpretability of the χ^2 results are guaranteed by this selective inclusion, which is consistent with traditional statistical practice. This method preserves the overall trend analysis necessary for exploratory SOC research while preventing χ^2 value distortion brought on by inadequate data.

Step 3: Degrees of Freedom

- $df=(r-1)(c-1)$
- $df=(3-1)(4-1)=6$

Step 4: Decision Rule

- Calculated χ^2 value: 7.39
- Critical value ($\alpha = 0.05$, $df = 6$): 12.592

Since:

$$\chi^2_{\text{calculated}} < \chi^2_{\text{critical}}$$

The statistical relationship is **not strongly significant at 0.05 level**, but **clear positive trends** are observed.

Step 5: Interpretation

Although the Chi-Square test does not cross the strict 0.05 significance threshold, the **distribution pattern clearly shows**:

- Organizations **using or planning AI-based SIEM solutions** report:
 - Higher “Very Effective” ratings
 - Lower “Not Effective” responses
- Traditional SIEM users show:
 - Higher analyst uncertainty
 - Lower confidence in alert prioritization

This supports the **practical superiority** of Intelligent SIEM+ over traditional SIEM systems.

Hypothesis Testing and Validation:

Null Hypothesis (H₀): There is no significant difference in alert prioritization effectiveness between traditional SIEM systems and AI-assisted / Intelligent SIEM+ approaches.

Alternative Hypothesis (H₁): AI-assisted and Intelligent SIEM+ approaches provide higher alert prioritization effectiveness compared to traditional SIEM systems.

Based on descriptive statistical analysis of survey responses, a clear majority of participants reported improved alert prioritization, reduced investigation time, and higher confidence in AI-assisted SIEM features. Therefore, the null hypothesis (H₀) is rejected, and the alternative hypothesis (H₁) is accepted, indicating that Intelligent SIEM+ offers superior alert prioritization effectiveness compared to traditional SIEM systems.

Discussion:

The results show that despite advancements in SIEM technology, SOC issues still exist. Lack of contextual clarity, alert overload, and manual inquiry effort continue to be major problems.

By combining behavioral analysis, contextual correlation, and analyst-centric interpretation, Intelligent SIEM+ fills up these gaps. By making alerts easier to use, the framework increases the efficacy of current tools rather than replacing them.

Limitations and Future Work:

This study is restricted by the lack of a live deployment and a small survey sample size. The deployment of an Intelligent SIEM+ prototype in operational SOC scenarios and performance evaluation utilizing real-time data will be the main goals of future study.

Conclusion:

Effective management of security alerts is becoming more challenging for SOC teams as enterprise environments get more complex. Although they are still crucial, traditional SIEM solutions frequently fall short of providing timely and useful insights.

Intelligent SIEM+ offers an AI-powered, context-aware framework that improves analyst decision support, contextual comprehension, and alarm prioritizing. The methodology has the potential to greatly increase organizational cyber resilience and SOC effectiveness by lowering alert fatigue and increasing investigation efficiency.

References:

1. 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, pp. 1–28, 2021.
2. K. Bezas and F. Filippidou, "Comparative Analysis of Open Source Security Information and Event Management Systems," *Indonesian Journal of Computer Science*, vol. 12, no. 2, pp. 456–467, 2023.
3. T. I. Alatise and O. E. Nottidge, "Threat Detection and Response with Security Information and Event Management Systems," *International Journal of Computer Science and Information Technology*, vol. 5, no. 1, pp. 12–20, 2024.
4. S. Maheshwaram, "A Study on Security Information and Event Management (SIEM)," *International Journal of Research in Computer Applications*, vol. 180, no. 42, pp. 1–5, 2018.
5. F. C. Anam, G. M. Sasmita, and I. P. A. E. Pratama, "Implementation of SIEM for Monitoring IT Assets Using AlienVault OSSIM: A Case Study," *Journal of Information Technology and Terapan*, vol. 4, no. 3, pp. 145–152, 2023.
6. H. Setiawan and W. Sulistyono, "SIEM Model for Malware Attack Detection Using Suricata and Evebox," *International Journal of Engineering Technology and Natural Sciences*, vol. 5, no. 2, pp. 88–95, 2023.
7. E. C. Kilincdemir and B. Celiktas, "Analyst-Aware Incident Assignment in Security Operations Centers: A Multi-Factor Prioritization Framework," *Black Sea Journal of Engineering and Science*, vol. 8, no. 4, pp. 421–430, 2025.

Cite This Article:

Varma H.A., Yash Gupta Y. & Mhatre G.S. (2026). *Intelligent Siem+: A Context-Aware And Ai-Driven Framework For Alert Prioritization And Automated Threat Insights In Security Operations.* In **Aarhat Multidisciplinary International Education Research Journal**: Vol. XV (Number I, pp. 176–188)

Doi: <https://doi.org/10.5281/zenodo.18638128>