

Explainable Analysis of Security Information and Event Management(SIEM) Event Logs for Event Profiling and Preliminary Triage

Pranisa Israsena

pranisa@tni.ac.th

Faculty of Information Technology

Thai-Nichi Institute of Technology, Bangkok, Thailand

*Amonpan Chomklin**

amonpan@tni.ac.th

Faculty of Information Technology

Thai-Nichi Institute of Technology, Bangkok, Thailand

Lalita Na Nongkhai

lalita@tni.ac.th

Faculty of Information Technology

Thai-Nichi Institute of Technology, Bangkok, Thailand

Received: February 16, 2026; Revised: June 20, 2026; Accepted: June 24, 2026; Published: June 30, 2026

ABSTRACT - Security Information and Event Management (SIEM) systems generate large volumes of heterogeneous event logs, imposing substantial investigation workloads on Security Operations Centers (SOCs). This study proposes an explainable framework for SIEM event analysis and preliminary triage that emphasizes interpretability and operational applicability without requiring labeled training data. The framework integrates descriptive and temporal analysis, structural association analysis using Spearman correlation, K-means clustering with TF-IDF features to construct interpretable event profiles, and explainable triage scoring based on standard SIEM indicators, including severity, priority, confidence, and aggregated counts. Experimental evaluation using operational SIEM logs (n=3,490 events) demonstrates 93.8% workload reduction while retaining 98.97% of high-severity events. External validation with labeled attack data (489 SQL injection attempts and 301 brute-force attacks) prioritized 98.2% of known attacks (776/790) with a 1.8% false-negative rate, outperforming severity-only (78.5%) and severity+priority (89.1%) baselines. The results confirm that unsupervised event profiling combined with explainable triage scoring can prioritize security-relevant events while filtering routine monitoring noise. The framework provides a practical, transparent analytical baseline for SOC environments under realistic operational constraints.

KEY WORDS -- SIEM, Security event analysis, Event clustering, Explainable analytics, Event triage, SOC operations

1. Introduction

Modern organizations increasingly rely on information systems and computer networks to support critical operations, which has led to a corresponding rise in cybersecurity risks arising from both external attacks and internal anomalies. Examples of such events include port scanning activities, repeated authentication attempts, unauthorized service access, as well as incidents caused by misconfigurations or software vulnerabilities. When these events occur, organizations must be capable of timely monitoring, detection, and response in order to mitigate potential impacts on systems and sensitive data. One of the widely adopted infrastructures for managing security-related information is the Security Information and Event Management (SIEM) system. SIEM systems collect event logs from diverse sources, such as network devices, security appliances, endpoints, and applications, and subsequently store, analyze, and visualize these data to support security monitoring and incident response [1]. Despite their ability to centralize security event data, SIEM systems face a fundamental operational challenge: the sheer volume and diversity of events generated on a daily basis, which significantly increases investigation workload and noise within the event stream [2]. A large proportion of events recorded in SIEM logs correspond to routine monitoring activities, such as recurring alerts from similar behavior patterns, automated scanning activities, or low-severity events that occur frequently. In contrast, events that are truly critical and require immediate response typically constitute only a small fraction of the total volume. This imbalance makes event prioritization and escalation highly time-consuming and heavily dependent on analyst expertise, and it directly contributes to the problem of alert fatigue in Security Operations Centers (SOCs) [2]. If event prioritization is not performed effectively, critical security incidents may be overlooked or response actions may be delayed. These challenges highlight that, while SIEM event logs are valuable for security monitoring, they require analytical approaches that can summarize global event characteristics and structure large volumes of data into interpretable forms. This need is particularly pronounced in organizational contexts where explainable and reusable analytical frameworks are preferred over opaque decision-making processes, consistent with established findings in explainable artificial intelligence and human-centered security analytics [3]. Common analytical approaches in security event analysis include descriptive analysis to summarize event trends and frequencies, structural association analysis of key indicators such as severity levels or correlated event counts, and clustering techniques to construct event profiles that reduce data

complexity to a manageable and interpretable level. In particular, the integration of structured indicators with text-derived representations, such as term frequency-inverse document frequency (TF-IDF), combined with clustering algorithms (e.g., k-means), has been widely adopted for summarizing large-scale security logs and related textual event records [5], [6].

Although machine learning techniques have been increasingly explored to automate event prioritization and detection, their deployment in real-world security operations remains constrained by several factors. These include the scarcity of high-quality labeled data, the inherent mismatch between closed-world learning assumptions and operational security environments, and the limited robustness of trained models when exposed to unseen behaviors [7]. Additional challenges include the limited explainability of many machine learning models, the presence of concept drift caused by evolving system configurations and attacker strategies, and the continuous maintenance burden associated with retraining and validating models in dynamic environments [3], [4]. Consequently, many organizations still require analytical and triage frameworks that can be applied immediately, remain interpretable to human analysts, and serve as a foundation for more advanced automation in later stages [2].

Motivated by these considerations, this study proposes a structured approach for analyzing SIEM event logs that begins with controlled data preparation to preserve the characteristics of real operational events. The methodology then applies descriptive and temporal analysis to explain event occurrence patterns, followed by structural association analysis to examine relationships among key SIEM indicators. Event clustering is subsequently employed to construct meaningful event profiles, and a preliminary triage framework is designed using existing SIEM indicators to compute interpretable triage scores. The framework is evaluated from an operational perspective, focusing on its ability to reduce investigation workload while retaining high-severity events.

The objectives of this study are threefold: (1) to analyze and summarize SIEM event log characteristics in temporal and structural dimensions, (2) to construct interpretable event profiles through clustering that integrates structural and textual features, and (3) to propose and evaluate an explainable preliminary triage framework in terms of investigation workload reduction and retention of high-severity events. The scope of this study focuses on SIEM event logs collected from a real operational

environment and emphasizes explainable analysis and preliminary triage rather than supervised prediction or prioritization performance evaluation. This scope is intended to establish a clear analytical baseline that can be systematically extended in future research.

The remainder of this paper is organized as follows. Section 2 reviews related work on SIEM analytics and event clustering. Section 3 describes the research methodology and data analysis procedures. Section 4 presents the experimental results, followed by discussion and practical implications in Section 5. Finally, Section 6 concludes the paper and outlines directions for future research.

2. Related Work

Security Information and Event Management (SIEM) systems constitute a fundamental infrastructure for monitoring and managing cybersecurity events by collecting and processing event logs from heterogeneous sources, including network devices, servers, and application systems. As discussed in Section 1, SIEM platforms enable centralized visibility over security-relevant activities; however, converting large volumes of heterogeneous and noisy log data into actionable decision support remains a major operational challenge [1], [2]. Standardization-oriented references, such as NIST Special Publication 800-92, emphasize best practices for log management—including time synchronization, secure storage, integrity preservation, and retention policies—to ensure that logs can be reliably used for incident investigation and forensic analysis [8]. While such standards establish an essential operational baseline, they do not prescribe data-driven analytical methods or systematic mechanisms for event summarization, prioritization, or triage.

From a data analytics perspective, clustering has been widely studied as a means to reduce the complexity of large-scale security logs and to identify recurring patterns. Landauer et al. provide a comprehensive survey of log clustering approaches in cybersecurity applications, covering distance-based, density-based, and hierarchical techniques applied to system and security logs [9]. Their survey highlights clustering as a practical unsupervised strategy for structuring massive log datasets in the absence of labels. Nevertheless, it also identifies critical limitations for operational adoption: clustering results are often difficult to interpret for analysts, and feature selection strongly influences the semantic meaning and usefulness of the resulting clusters [9]. These findings suggest that clustering alone is insufficient

for supporting analyst workflows unless complemented by interpretable summaries or profiles.

Several applied studies have explored machine learning techniques for SIEM log analysis, particularly the integration of unstructured textual log fields with structured indicators. A common approach involves transforming textual attributes (e.g., event names or descriptions) into vector representations using term frequency–inverse document frequency (TF-IDF) weighting [6], followed by clustering algorithms such as K-means [12]. Tendikov et al. demonstrate that such SIEM-oriented analytics pipelines can efficiently group events with similar patterns at relatively low computational cost, making them attractive as baseline analytical methods in operational environments [5]. However, these studies primarily focus on algorithmic feasibility or grouping performance and pay limited attention to how the resulting clusters can be operationalized as interpretable event profiles that directly support investigation and triage activities in Security Operations Centers (SOCs) [5], [9].

Beyond event grouping, alert prioritization and triage have emerged as critical research areas in response to alert fatigue and excessive investigation workloads, as highlighted in Section 1 [2]. Wang et al. propose AlertPro, a context-aware alert prioritization framework that leverages alert sequences, contextual signals, and feedback mechanisms to improve triage efficiency in SOC [10]. Complementing this line of work, Jalalvand et al. provide a systematic survey of alert prioritization criteria and methods, revealing substantial diversity in prioritization strategies and emphasizing the strong influence of data availability, contextual enrichment, and organizational constraints on practical deployment [11]. While these approaches demonstrate the potential benefits of prioritization, they often assume the availability of rich contextual or historical information that may not be consistently present across different SIEM deployments, thereby limiting their general applicability.

To avoid citation inconsistency and clarify the methodological positioning of this study, Table 1 summarizes representative prior work using only studies that are included in the reference list. The comparison focuses on log management, SIEM analytics, log clustering, alert prioritization, and explainable security analytics. These studies collectively show that existing approaches often address event collection, clustering, prioritization, or explainability as separate components, whereas this study integrates unsupervised event profiling with an explainable triage scoring mechanism using commonly available SIEM indicators.

Table 1. Summary of related work in SIEM analytics and security event triage

Study	Main focus	Method / Approach	Key contribution	Limitation / Gap related to this study
González-Granadillo et al. [1]	SIEM systems	SIEM analysis and trend review	Summarizes SIEM usage, trends, and challenges in security monitoring	Does not provide a lightweight event profiling and triage framework
NIST SP 800-92 [8]	Log management	Operational guideline	Provides best practices for log collection, storage, integrity, and retention	Focuses on log management practice rather than data-driven event profiling
Landauer et al. [9]	Log clustering	Survey of clustering approaches	Shows that clustering is useful for reducing log complexity in cybersecurity	Highlights interpretability and feature selection challenges
Tendikov et al. [5]	SIEM analytics	Machine-learning-based SIEM data acquisition and analysis	Demonstrates practical analytics pipelines for SIEM data	Limited connection between clustering outputs and SOC triage decisions
Wang et al. [10]	Alert prioritization	Context-aware alert prioritization	Addresses alert fatigue through prioritization of security alerts	Requires richer contextual and historical information
Jalalvand et al. [11]	SOC alert prioritization	Systematic survey	Summarizes prioritization criteria and methods in SOC environments	Shows lack of a unified lightweight baseline for heterogeneous SIEM settings
This study	Explainable SIEM event profiling and triage	TF-IDF + K-means event profiling + fixed-weight triage scoring	Integrates unsupervised event profiles with explainable preliminary triage	Limited to a single-day primary dataset and two labeled attack categories

As summarized in Table 1, prior studies provide important foundations for SIEM monitoring, log management, clustering, and alert prioritization. However, many approaches either focus on operational log management, algorithmic clustering, or context-heavy prioritization. This study addresses the gap by proposing a lightweight and explainable framework that connects unsupervised event profiling with preliminary triage scoring without requiring labeled training data or complex contextual enrichment.

3. Methodology

This study aims to develop a structured and explainable approach for analyzing Security Information and Event Management (SIEM) event logs in order to support the understanding of security event characteristics and preliminary event triage in real operational settings. The methodological design emphasizes interpretability, reliance on commonly available SIEM indicators, and systematic reproducibility, enabling the proposed approach to be applied and extended across diverse organizational contexts

3.1 Conceptual Research Framework

The conceptual framework is organized as a sequential pipeline from raw SIEM event logs to decision-support outputs. It consists of data preparation, descriptive analysis, exploratory analysis, structural association analysis, event profiling via clustering, preliminary triage scoring, and validation with labeled attack data. Descriptive analysis summarizes event volume and temporal distribution, whereas exploratory analysis examines contextual event composition, including event categories, source countries, destination ports, device or product types, and textual characteristics. Raw SIEM logs are first converted into an analytical dataset through timestamp parsing, data type conversion, missing-value handling, preservation of potentially anomalous values, and consolidation of textual fields. The transformed data are then represented using standardized SIEM indicators and TF-IDF textual features for clustering and triage scoring. This framework is intended to support SOC analysts by transforming large volumes of raw events into interpretable profiles and prioritized investigation outputs, rather than replacing analyst judgment.




Figure 1. Conceptual framework and data construction pipeline for explainable SIEM event profiling and preliminary triage.

As shown in Figure 1, each analytical stage produces outputs that support the next stage, enabling transparent decision support for SOC analysts. The framework transitions from data preparation to event profiling and triage scoring, while remaining grounded in standard SIEM indicators and interpretable analytical outputs.

3.2 Research Questions and Hypotheses

To guide the analysis in a systematic and reproducible manner, this study defines research questions and hypotheses that are examined using empirical evidence, with an emphasis on structural characteristics and operational applicability rather than predictive performance. This methodological choice reflects the absence of ground-truth event labels in the primary SIEM dataset and aligns with the exploratory and descriptive nature of the proposed framework. The research questions and hypotheses are explicitly aligned with the sequential stages of the conceptual framework described in Section 3.1, encompassing data characterization, structural association analysis, event profiling via clustering, evaluation of the preliminary triage framework, and validation with labeled attack data. To enhance clarity, traceability, and consistency between analytical objectives and methodological design, the research questions and their corresponding hypotheses are summarized in Table 2, together with their analytical focus and associated framework stages.

Table 2. Research questions and hypotheses aligned with the conceptual framework

RQ/Hypothesis	Research Question	Hypothesis	Framework Stage
RQ1 / H1	What are the temporal characteristics and overall patterns of SIEM event logs?	SIEM event logs exhibit interpretable temporal patterns with observable temporal concentrations and event composition characteristics.	Descriptive and temporal analysis
RQ2 / H2	What structural associations exist among commonly used SIEM indicators?	Event severity levels and correlated or aggregated event counts exhibit meaningful associations from an operational perspective.	Structural association analysis
RQ3 / H3	Can event clustering produce interpretable event profiles with practical significance?	Integration of structural and textual SIEM features enables the construction of distinct and interpretable event profiles with practical significance.	Event profiling via clustering
RQ4 / H4	Does the preliminary triage framework achieve workload reduction while retaining high-severity events?	An explainable triage framework based on standard SIEM indicators achieves a favorable trade-off between efficiency and risk.	Preliminary triage evaluation
RQ5 / H5	Can the framework effectively prioritize genuine security incidents when validated with labeled attack data?	The framework achieves high prioritization rates for labeled attacks while maintaining workload reduction, demonstrating effective differentiation between genuine security incidents and routine events.	Validation with labeled attack data

3.3 Data Source and Scope of the Study

The primary dataset consists of 3,490 SIEM event logs collected from a real operational environment within

a single-day observation window. The dataset contains structured SIEM indicators, including severity, priority, confidence or relevance score, device or product information, event categories, and

correlated or aggregated event counts, as well as unstructured textual fields such as event names and descriptions. The primary dataset is unlabeled and is therefore used for descriptive analysis, exploratory analysis, structural association analysis, unsupervised event profiling, and preliminary triage evaluation. The raw SIEM logs were converted into an analytical dataset through a controlled preparation pipeline. This included timestamp parsing, numerical type conversion, missing-value handling, preservation of potentially anomalous values, and consolidation of textual fields for TF-IDF feature extraction. Descriptive analysis was then used to summarize event volume and temporal distribution, whereas exploratory analysis examined event composition, including behavioral categories, source countries, destination ports, device or product types, and recurring textual characteristics.

Supplementary labeled attack data, consisting of SQL injection attempts and brute-force authentication attacks, were used only for external validation of the triage framework. These labeled events were not used to train the clustering model, optimize triage weights, or adjust the triage threshold. This design preserves the label-free nature of the proposed framework while allowing its prioritization capability to be assessed under realistic SOC conditions.

3.4 Structural and Association Analysis

Given the skewed distributions and the presence of outliers commonly observed in SIEM event data, this study employs non-parametric statistical analysis to examine structural relationships among key SIEM indicators. In particular, associations between event severity levels and correlated or aggregated event counts are analyzed using Spearman's rank correlation coefficient (ρ), which is defined as:

$$\rho = 1 - (6\sum d_i^2) / (n(n^2 - 1)) \quad (1)$$

where d_i denotes the difference between the ranks of paired observations for event i , and n represents the total number of events. Spearman's rank correlation is suitable for ordinal variables such as severity levels and does not require assumptions of normality or linearity, making it appropriate for analyzing SIEM data characterized by non-Gaussian distributions and extreme values [13].

The objective of this analysis is not to infer causal relationships, but to provide empirical justification for the selection of structural indicators used in subsequent analytical stages. Specifically, the observed associations support data-driven decisions regarding which indicators are integrated into event clustering and preliminary triage scoring. In this sense, the structural association analysis serves as a

methodological bridge between descriptive characterization of SIEM events and the construction of interpretable event profiles in later stages of the framework.

3.5 Event Profiling via Clustering

Event clustering was used to construct interpretable event profiles by combining structured SIEM indicators with textual features extracted from event names and descriptions. Textual fields were represented using TF-IDF [6], while structured indicators included severity, priority, confidence or relevance score, and correlated or aggregated event count. Numerical indicators were standardized before concatenation with TF-IDF features to reduce scale bias. The number of clusters k was determined using the silhouette coefficient, which provides an internal validation measure of clustering structure and supports data-driven selection of an appropriate cluster granularity. For an event i , the silhouette coefficient $s(i)$ is defined as:

$$s(i) = (b(i) - a(i)) / \max\{a(i), b(i)\} \quad (2)$$

where $a(i)$ denotes the average distance between event i and all other events within the same cluster, and $b(i)$ represents the minimum average distance between event i and events in the nearest neighboring cluster. Values of $s(i)$ range from -1 to 1 , with higher values indicating better separation between clusters [14].

K-means was adopted as a transparent baseline clustering method [12]. The number of clusters was selected using silhouette-guided assessment [14] together with interpretability considerations. The final configuration used $k = 10$, producing ten event profiles labeled Cluster 0 to Cluster 9. A fixed random seed and multiple centroid initializations were used to improve reproducibility and reduce sensitivity to initialization.

The resulting clusters were interpreted as event profiles rather than attack-type labels. Each profile was summarized using cluster size, dominant behavior, dominant device or product, average priority, average confidence, and representative textual characteristics. These profiles serve as interpretable abstractions of recurring SIEM event patterns and provide inputs to the preliminary triage framework.

3.6 Preliminary Triage Framework and Evaluation

The preliminary triage framework is designed to reduce investigation workload while retaining critical security events, in alignment with the operational

objectives of Security Operations Centers (SOCs). To this end, a triage score is computed using standard SIEM indicators, without optimizing parameter weights, in order to preserve transparency and interpretability. For each event i , the triage score is defined as:

$$\text{Triage Score}_i = w_1 \cdot S_i + w_2 \cdot P_i + w_3 \cdot C_i + w_4 \cdot \log(1 + A_i) \quad (3)$$

In Equation (3), S_i denotes the event severity level, P_i represents the event priority, C_i corresponds to a normalized confidence or relevance score provided by the SIEM system, and A_i is the number of correlated or aggregated events associated with event i . Reported immediately here for structural transparency, the fixed weights are defined as $w_1 = 1.0$, $w_2 = 0.8$, $w_3 = 0.5$, and $w_4 = 0.3$. These values were not learned from the dataset and were not optimized to maximize detection performance. Instead, they were defined as an analyst-oriented prioritization policy to preserve transparency, reproducibility, and operational interpretability. Aggregated event count receives the lowest direct weight because frequent correlation can indicate coordinated or repeated activity, but raw count values may be highly skewed in SIEM environments.

A logarithmic transformation, $\log(1 + A_i)$, was therefore applied to the aggregated event count term to reduce the influence of extreme values while preserving monotonicity. This design allows events with larger correlated counts to contribute meaningfully to the triage score without allowing high-volume routine events to dominate the prioritization process. The scoring framework is therefore intended as an explainable triage heuristic rather than a learned predictive model. This fixed-weight design avoids dataset-specific parameter tuning and supports direct interpretation of the triage score. For example, an event receives a higher score when it has higher severity, higher operational priority, stronger confidence, or a larger number of correlated events. Because each component is based on commonly available SIEM indicators, the resulting score can be inspected and explained to analysts without requiring complex model interpretation.

Evaluation of the preliminary triage framework focuses on two operationally relevant criteria. First, workload reduction is measured by the proportion of events excluded from manual investigation when applying the triage score threshold. Second, retention of critical events is assessed by examining the presence and proportion of high-severity events within the prioritized subset. Together, these criteria reflect the trade-off between efficiency and risk that SOC analysts routinely manage. To enhance reliability and traceability, all intermediate datasets, analytical outputs, and evaluation results are preserved in reusable file formats. Algorithmic parameters, scoring weights, and initialization settings are explicitly documented, enabling systematic reproduction and

comparative analysis in future studies. This approach promotes transparency and repeatability while avoiding reliance on dataset-specific parameter optimization.

3.6.1 Validation Procedure with Labeled Attack Data

To evaluate the triage framework's prioritization capability, the supplementary labeled attack data were processed through the same analytical pipeline used for the primary SIEM events. The validation data were used only for external assessment and were not used to retrain the clustering model, update cluster centroids, optimize triage weights, or adjust the triage threshold. The validation procedure consisted of three steps. First, labeled SQL injection and brute-force authentication events were transformed into the same analytical schema as the primary SIEM dataset, including severity, priority, confidence, textual descriptions, and aggregated event counts. Second, the transformed events were assigned to the existing ten event profiles generated by the primary K-means model without retraining or updating centroids. Third, triage scores were computed using Equation (3) with the same fixed weights defined in Section 3.6.

Validation performance was assessed using attack prioritization rate, false-negative rate, and category-specific prioritization rate. This procedure reflects a realistic SOC scenario in which newly observed events are mapped to existing event profiles and prioritized using a transparent triage policy rather than a retrained supervised classification model.

4. Results

The experimental results are presented in accordance with the defined research questions (RQ1–RQ5) and their corresponding hypotheses, in order to demonstrate consistency between the conceptual framework, the methodological design described in Section 3, and the empirical evidence obtained from real SIEM event logs. Each subsection reports results derived directly from the analytical procedures outlined in the Methodology and is supported by figures and tables generated from the processed dataset.

4.1 Temporal Characteristics and Overall Event Patterns (RQ1 / H1)

To address RQ1, the analysis begins with an examination of the temporal characteristics of SIEM event logs, following the descriptive and temporal analysis stage. This stage aims to establish a global understanding of event occurrence patterns within the available observation window and to assess whether the data exhibit interpretable temporal structures

suitable for further analysis. Figure 2 illustrates the daily event volume recorded by the SIEM system during the observed period. As shown in the figure, a total of 3,490 events were recorded on January 5, 2026, indicating that the system continuously generated security-related events at a high frequency throughout the day. Because the dataset represents a single-day snapshot, the analysis in this subsection focuses on descriptive characterization of event volume within the available temporal scope, rather than on identifying long-term trends or seasonal effects. This scope is consistent with the data source constraints described in Section 3.3 and avoids overinterpretation beyond the temporal limits of the dataset.




Figure 2. Daily event volume from SIEM logs (single-day snapshot).

To further examine intraday temporal patterns, event frequencies were analyzed across weekdays and hours, as described in the temporal analysis procedure. Figure 3 presents a heatmap of event frequency by weekday and hour, illustrating the distribution of events within the observed day. Although the temporal scope is limited, the results reveal non-uniform event distributions, with events concentrated during specific hours rather than evenly spread across the day. Such temporal concentration reflects typical SIEM operational behavior, where event generation is influenced by system usage patterns, automated monitoring processes, and scheduled activities within the monitored environment.

These observations support Hypothesis H1, which posits that SIEM event logs exhibit interpretable temporal patterns with observable temporal concentrations. Importantly, the identified temporal structures provide contextual grounding for subsequent analytical stages, including structural association analysis (Section 4.2) and event profiling via clustering (Section 4.3). By confirming the presence of meaningful temporal characteristics at the descriptive level, this subsection establishes that the dataset is suitable for deeper structural and profiling

analyses, as envisioned in the conceptual framework (Figure 1).




Figure 3. Heatmap of event frequency by day and hour.

4.2 Structural Associations Among SIEM Indicators (RQ2 / H2)

To address RQ2, this subsection examines structural associations among commonly used SIEM indicators, following the non-parametric association analysis. Specifically, the relationship between event severity and the number of correlated or aggregated events is analyzed using Spearman's rank correlation coefficient, which is suitable for ordinal variables and skewed SIEM data distributions.

The Spearman correlation analysis showed a positive association between event severity and correlated or aggregated event counts ($n = 3,490$). This result indicates that higher-severity events tended to be associated with a larger number of related events. The correlation result is not interpreted as evidence of causality; rather, it provides empirical justification for including aggregated event count together with severity in the subsequent clustering and preliminary triage scoring stages.

4.3 Event Clustering and Event Profile Construction (RQ3 / H3)

To address RQ3, this subsection reports the results of event clustering performed using the integration of structural indicators and textual features. The objective of this analysis is not to optimize clustering performance, but to construct interpretable event profiles that summarize recurring SIEM event patterns in a form suitable for operational use. The clustering results show that SIEM events can be partitioned into multiple groups with distinct structural and behavioral characteristics.

The distribution of events across clusters is illustrated in Figure 4, which reveals a highly imbalanced cluster size distribution. A small number

of clusters contain the majority of events, while several clusters are substantially smaller. This pattern is characteristic of real-world SIEM environments, where routine monitoring activities generate large volumes of repetitive, low-severity events, whereas potentially significant or specialized events occur less frequently. Specifically, as shown in Figure 4, the largest cluster (Cluster 0) dominates the dataset, indicating a concentration of common and repetitive event types. Additional medium-sized clusters (e.g., Clusters 2 and 6) capture secondary patterns of recurring behavior, while multiple small clusters

represent rare or distinctive event structures. This separation suggests that clustering effectively differentiates between high-frequency operational noise and lower-frequency events with potentially higher analytical relevance. To facilitate operational interpretation, the characteristics of each cluster are summarized as event profiles, presented in Table 3. Each profile aggregates key attributes of the corresponding cluster, including the number of events, dominant behavioral category, dominant device or product, average priority, and average model confidence.

Table 3. Summary of event profiles derived from clustering analysis.

Cluster	Number of Events	Dominant Behavior	Dominant Device/Product	Avg. Priority	Avg. Model Confidence
0	1850	Unknown	OBS	3.00	0.00
2	555	Unknown	OBS	3.00	4.00
6	426	/Access/Stop	Unix	3.02	0.00
7	264	Unknown	Threat inSights	3.00	0.00
1	195	Unknown	NAXSI	7.00	3.96
3	137	/Communicate/Query	Apache	3.00	0.00
9	49	Unknown	WebServer	2.61	0.00
4	23	Unknown	ArcSight	4.52	0.00
8	2	/Execute/Response	ArcSight	5.00	0.00
5	1	/Execute/Response	ArcSight	5.00	0.00

As shown in Table 3, clusters with large event counts (e.g., Clusters 0 and 2) are dominated by generic or “Unknown” behaviors with moderate average priority, reflecting routine SIEM monitoring noise. In contrast, smaller clusters (e.g., Clusters 1, 4, 8, and 5) exhibit higher average priority values and are associated with specific devices or products, such as NAXSI or ArcSight. Although these clusters contain relatively few events, their distinct structural characteristics highlight their potential operational relevance, particularly in the context of event prioritization. Overall, the combined evidence from Figure 4 and Table 3 demonstrates that event clustering effectively reduces the complexity of large-scale SIEM data by organizing events into a manageable number of interpretable profiles. These profiles provide structured representations of event behaviors that support analyst sense-making and directly inform the preliminary triage framework presented in Section 4.4. Accordingly, the results in this subsection support Hypothesis H3, confirming that the integration of structural and textual SIEM features enables the construction of distinct and interpretable event profiles with practical significance.




Figure 4. Distribution of events across clusters.

4.4 Evaluation of the Preliminary Triage Framework (RQ4 / H4)

To address RQ4, the preliminary triage framework is evaluated using the operational criteria defined in Section 3.6, namely investigation workload reduction and retention of critical events. The evaluation applies the triage score to rank all SIEM events and selects a prioritized subset for analyst investigation. Figure 5 illustrates the effect of the triage mechanism on investigation workload. By selecting the top 216 events based on the triage score—corresponding to 6.17% of the total event volume—the framework excludes approximately 93.83% of events from manual investigation. This substantial reduction demonstrates that the proposed indicator-based

scoring mechanism effectively filters routine and low-priority events, yielding a manageable investigation set without requiring supervised learning or dataset-specific parameter optimization.

In addition to workload reduction, the framework is evaluated in terms of its ability to retain critical events. As shown in Figure 6, of the 195 high-severity events present in the dataset, 193 events (98.97%) are retained within the shortlisted subset after triage. This result indicates that the aggressive reduction in investigation volume does not disproportionately exclude high-severity events.




Figure 5. Effect of the triage mechanism on investigation workload.

Taken together, these results support Hypothesis H4, confirming that an explainable triage framework based on standard SIEM indicators can achieve a favorable trade-off between efficiency and risk. Importantly, the framework operates without supervised learning, dataset-specific parameter optimization, or complex contextual enrichment, aligning with the design objectives outlined in Section 3.6. The high retention rate of critical events, combined with substantial workload reduction, demonstrates the practical utility of the proposed approach as a baseline decision-support mechanism in operational SIEM environments.




Figure 6. Retention of high-severity events under the preliminary triage framework.

4.5 Validation with Labeled Attack Data (RQ5 / H5)

To externally validate the triage framework, supplementary labeled attack data were processed using the same analytical pipeline as the primary SIEM dataset. The labeled data included 489 SQL injection attempts and 301 brute-force authentication attacks. These events were used only for validation and were not used to train the clustering model, update centroids, optimize weights, or adjust the triage threshold.

The transformed attack events were assigned to the existing ten event profiles, Cluster 0 to Cluster 9, generated by the primary K-means model. This preserves consistency between the unsupervised event profiling stage and the validation analysis. Brute-force authentication events were mainly associated with repeated authentication patterns, whereas SQL injection attempts were distributed across several existing event profiles due to their more diverse textual characteristics. Figure 7 shows the distribution of labeled attack events across the existing ten event profiles.




Figure 7. Distribution of labeled attack events across the existing ten event profiles generated by the primary K-means model.

Triage scores were computed for all labeled attack events using Equation (3) with the same fixed weights used in the primary triage evaluation. Table 4 summarizes the prioritization performance of the preliminary triage framework on the labeled attack validation data.

The framework prioritized 776 of 790 labeled attack events, yielding an overall attack prioritization rate of 98.2% and a false-negative rate of 1.8%. SQL injection attempts were prioritized at 97.1% (475/489), while brute-force authentication attacks were prioritized at 100.0% (301/301). All 14 false negatives were SQL injection attempts with relatively low aggregated event counts and triage scores just below the threshold. Figure 8 summarizes the prioritization performance across attack categories.

Table 4. Triage framework prioritization performance on labeled attack validation set

Attack Category	Total Events	Prioritized	Prioritization Rate	False Negatives
SQL Injection	489	475	97.1%	14 (2.9%)
Brute-force Authentication	301	301	100.0%	0 (0.0%)
Overall	790	776	98.2%	14 (1.8%)

To assess the contribution of the integrated triage framework beyond simpler prioritization heuristics, Table 5 compares prioritization performance across different prioritization strategies applied to the labeled attack validation set.




Figure 8. Prioritization performance summary across labeled attack categories.

Table 5 shows that the proposed framework outperformed random selection, severity-only, and severity-plus-priority baselines under the same 6.17%

workload allocation. The proposed framework achieved 98.2% attack prioritization with a 1.8% false-negative rate, compared with 78.5% for severity-only and 89.1% for severity-plus-priority. This result indicates that combining severity, priority, confidence, aggregated event count, and clustering-derived profiles provides stronger prioritization performance than simpler heuristics.

Table 5. Comparison of prioritization strategies on labeled attacks

Prioritization Strategy	Workload	Attack Prioritization Rate	False Negative Rate
Random Selection	6.17%	6.2% $\pm$ 2.1%	93.8%
Severity Only	6.17%	78.5%	21.5%
Severity + Priority	6.17%	89.1%	10.9%
Proposed Framework (Full)	6.17%	98.2%	1.8%

Note: All strategies assume consistent 6.17% workload allocation. Random selection prioritization rate represents mean $\pm$ standard deviation from 100 iterations. The proposed framework integrates severity, priority, confidence, aggregated event counts, and clustering-based profiles

The validation analysis using labeled attack data confirms that the proposed triage framework achieves a favorable balance between workload reduction and prioritization performance. With a 98.2% attack prioritization rate for known attacks and only 1.8% false negatives, the framework demonstrates practical viability for operational deployment in SOC environments. The integration of event clustering, structural indicators, and explainable triage scoring enables effective prioritization of genuine security incidents while filtering routine monitoring noise, without requiring supervised learning, extensive labeling infrastructure, or complex contextual enrichment.

The observed false negatives (n=14, all SQL injection attempts) highlight specific operational scenarios where the current framework exhibits limitations: initial reconnaissance attempts with low correlation counts before sustained attack patterns emerge, and events receiving moderate severity assignments from source systems. These findings suggest potential avenues for future enhancement, such as incorporating rare event detection heuristics or adaptive threshold mechanisms that account for temporal attack evolution. Nonetheless, the validation results provide strong empirical support for the framework's core objective: supporting analyst

decision-making through transparent, interpretable, and effective event triage under realistic operational constraints.

To consolidate the alignment between the research questions, the conceptual framework, and the empirical findings, Table 6 presents a structured mapping of the analytical stages and their corresponding supporting evidence. This table provides a high-level synthesis of the results reported in Sections 4.1–4.5 and is intended to emphasize traceability rather than to introduce new experimental findings. As summarized in Table 6, each research question is addressed through a clearly defined analytical stage within the conceptual framework, and each stage is supported by empirical evidence derived from real SIEM log data. This mapping demonstrates the internal consistency of the proposed methodology and confirms that the analytical pipeline—from descriptive characterization to preliminary triage—operates as intended. Collectively, the results establish a coherent empirical foundation for the discussion of implications, limitations, and future research directions presented in the subsequent section.

Table 6. Mapping between research questions, framework stages, and supporting evidence.

Research Question	Framework Stage	Supporting Evidence
RQ1	Descriptive & Temporal Analysis	Figure 2–3: 3,490 events (single-day snapshot) with interpretable temporal concentration and event composition patterns
RQ2	Structural & Association Analysis	Positive Spearman association ($n = 3,490$).
RQ3	Event Profiling via Clustering	Table 3 and Figure 4: Distinct and interpretable event profiles with operational significance
RQ4	Preliminary Triage Evaluation	Figure 5–6: 93.8% workload reduction with 98.97% retention of high-severity events
RQ5	Validation with Labeled Attack Data	Table 4–5, Figure 7–8: 98.2% attack prioritization rate for

Research Question	Framework Stage	Supporting Evidence
		labeled attacks (776/790) with 1.8% false negative rate

5. Discussion

The results demonstrate that the proposed framework can reduce SIEM investigation workload while preserving analyst-relevant security events. Descriptive and exploratory analysis characterize the operational event environment, while clustering converts heterogeneous logs into interpretable event profiles. The triage score then provides a transparent prioritization mechanism based on severity, priority, confidence, and aggregated event count. This design is suitable for SOC settings where labeled training data may be limited but explainable decision support is required.

Although the primary operational dataset is unlabeled and limited to a single-day snapshot, supplementary labeled attack data were used for external validation of prioritization performance. The validation results indicate that the framework can prioritize most labeled attack events without retraining the clustering model or tuning triage parameters. Nevertheless, the study remains limited by the short observation window and validation using only two attack categories. Future work should evaluate longer-term SIEM datasets, broader attack types, and analyst-in-the-loop deployment.

6. Conclusion

This study proposed an explainable framework for SIEM event profiling and preliminary triage using commonly available SIEM indicators and textual event information. The framework integrates data preparation, descriptive and exploratory analysis, structural association analysis, K-means-based event profiling, and fixed-weight triage scoring. Its main contribution is a transparent and label-free analytical baseline that supports early-stage event prioritization without requiring supervised model training.

Experimental results using 3,490 operational SIEM events showed that the proposed triage framework reduced the investigation workload by 93.8% while retaining 98.97% of high-severity events. External validation using labeled attack data further showed that the framework prioritized 98.2% of known attacks, including SQL injection and brute-force authentication events, without using these labeled data for training or parameter optimization.

Overall, the results indicate that unsupervised event profiling combined with explainable triage scoring can provide practical decision support for SOC analysts under realistic operational constraints. The main limitations are the single-day primary dataset and validation using only two attack categories. Future work should evaluate longer observation periods, broader attack types, and analyst-in-the-loop deployment while preserving the framework's interpretability.

Acknowledgments

The authors would like to thank So Secure Company Limited for providing access to anonymized SIEM event log data used in this study. The dataset was collected from a real operational environment and was made available solely for academic research purposes. All data were handled in accordance with applicable confidentiality and privacy requirements, and no information that could identify specific systems, organizations, or individuals is disclosed in this paper. The support from So Secure Company Limited was instrumental in enabling the empirical evaluation of the proposed analytical framework under realistic operational conditions.

References

- [1] G. González-Granadillo, S. González-Zarzosa, and R. Diaz, "Security Information and Event Management (SIEM): Analysis, trends, and usage in critical infrastructures," *Sensors*, vol. 21, no. 14, Art. no. 4759, 2021, doi: 10.3390/s21144759.
- [2] S. Tariq, M. B. Chhetri, S. Nepal, and C. Paris, "Alert fatigue in security operations centres: Research challenges and opportunities," *ACM Computing Surveys*, vol. 57, no. 9, Art. no. 224, 2025, doi: 10.1145/3723158.
- [3] R. Guidotti et al., "A survey of methods for explaining black box models," *ACM Computing Surveys*, vol. 51, no. 5, Art. no. 93, 2018, doi: 10.1145/3236009.
- [4] J. Gama, I. Žliobaitė, A. Bifet, M. Pechenizkiy, and A. Bouchachia, "A survey on concept drift adaptation," *ACM Computing Surveys*, vol. 46, no. 4, Art. no. 44, 2014, doi: 10.1145/2523813.
- [5] N. Tendikov et al., "Security Information Event Management data acquisition and analysis methods with machine learning principles," *Results in Engineering*, vol. 22, Art. no. 102254, 2024, doi: 10.1016/j.rineng.2024.102254.
- [6] G. Salton and C. Buckley, "Term-weighting approaches in automatic text retrieval," *Information Processing & Management*, vol. 24, no. 5, pp. 513–523, 1988, doi: 10.1016/0306-4573(88)90021-0.
- [7] R. Sommer and V. Paxson, "Outside the closed world: On using machine learning for network intrusion detection," in *Proc. IEEE Symposium on Security and Privacy*, 2010, pp. 305–316, doi: 10.1109/SP.2010.25.
- [8] K. Scarfone and M. Souppaya, *Guide to Computer Security Log Management*, NIST Special Publication 800-92, National Institute of Standards and Technology, 2006, doi: 10.6028/NIST.SP.800-92.
- [9] M. Landauer, F. Skopik, M. Wurzenberger, and A. Rauber, "System log clustering approaches for cyber security applications: A survey," *Computers & Security*, vol. 92, Art. no. 101739, 2020, doi: 10.1016/j.cose.2020.101739.
- [10] X. Wang et al., "Combating alert fatigue with AlertPro: Context-aware alert prioritization using reinforcement learning for multi-step attack detection," *Computers & Security*, vol. 137, Art. no. 103583, 2024, doi: 10.1016/j.cose.2023.103583.
- [11] F. Jalalvand, M. B. Chhetri, S. Nepal, and C. Paris, "Alert prioritisation in security operations centres: A systematic survey on criteria and methods," *ACM Computing Surveys*, vol. 57, no. 2, Art. no. 42, 2025, doi: 10.1145/3695462.
- [12] S. P. Lloyd, "Least squares quantization in PCM," *IEEE Transactions on Information Theory*, vol. 28, no. 2, pp. 129–137, 1982, doi: 10.1109/TIT.1982.1056489.
- [13] C. Spearman, "The proof and measurement of association between two things," *The American Journal of Psychology*, vol. 15, no. 1, pp. 72–101, 1904, doi: 10.2307/1412159.
- [14] P. J. Rousseeuw, "Silhouettes: A graphical aid to the interpretation and validation of cluster analysis," *Journal of Computational and Applied Mathematics*, vol. 20, pp. 53–65, 1987, doi: 10.1016/0377-0427(87)90125-7.