International Journal of Technology Management & Humanities (IJTMH)
e-ISSN: 2454 – 566X, Volume 06, Issue 1-2, (January-June 2020), www.ijtmh.com

A Risk-Centric Methodology for Detecting Cyber Threats Across SAP Landscapes Using Security Event Correlation, Statistical Anomaly Profiling, and Enterprise-Wide SIEM Analytics

Rajendra Asuri
Senior SAP integration consultant

Abstract
This study presents a risk-centric methodology for detecting cyber threats across SAP landscapes through security event correlation, statistical anomaly profiling, and enterprise-wide Security Information and Event Management (SIEM) analytics. The research addresses the limitations of fragmented log monitoring, isolated security alerts, and conventional rule-based detection in complex SAP environments. A mixed-methods approach combines qualitative analysis of SAP security events and attack scenarios with quantitative evaluation of event frequency, behavioral deviations, alert correlation, and detection accuracy. The proposed methodology integrates SAP audit logs, application events, authentication records, and system activity indicators into a centralized analytical framework for identifying suspicious behavior and prioritizing security incidents according to operational risk. Evaluation criteria include detection precision, false-positive rates, incident identification time, and correlation effectiveness. The framework establishes a structured relationship between user behavior, technical anomalies, business process exposure, and potential security violations, enabling investigators to distinguish routine operational variations from meaningful threat indicators. Its principal contribution lies in connecting SAP-specific security context with enterprise SIEM capabilities to improve investigative consistency, threat visibility, and incident response prioritization. The findings provide a foundation for more reliable security monitoring, improved audit readiness, and evidence-based cybersecurity decisions across interconnected enterprise systems. The methodology also contributes to academic research by offering a repeatable model for evaluating log-based threat detection and anomaly analysis. Its practical significance extends to SAP security teams, auditors, and enterprise risk managers seeking measurable improvements in threat identification, investigation efficiency, and protection of critical business information.

Keywords: SAP Security, Threat Detection, SAP Enterprise Threat Detection, SIEM Integration, Log Analysis, Anomaly Detection, Security Monitoring, Security Event Correlation, Incident Response, Audit and Compliance, Behavioral Analytics, Statistical Anomaly Profiling, Security Information and Event Management, Cyber Threat Intelligence, Intrusion Detection, Risk-Based Security Assessment, Security Event Management, User Activity Monitoring, Log Correlation, False-Positive Reduction, Enterprise Cybersecurity

DOI: 10.21590/ijtmh.06.1-2.03

1. Introduction

Enterprise Resource Planning (ERP) systems have become essential to the management of business operations, financial transactions, procurement, human resources, supply chains, and organizational information. SAP landscapes frequently support interconnected business processes in which the

confidentiality, integrity, and availability of information directly influence operational continuity. As enterprises expand their SAP environments across multiple application servers, databases, interfaces, and infrastructure components, the complexity of security monitoring increases considerably. These interconnected systems generate substantial volumes of application logs, audit records, authentication events, and transaction data. Examining these records systematically is necessary to identify unauthorized activities, suspicious access patterns, and potential security violations before they develop into significant business risks.

Conventional security monitoring approaches commonly rely on predefined rules, individual log reviews, and alerts generated by separate infrastructure components. Although these techniques remain useful for identifying known attack patterns, they may be insufficient when malicious activity is distributed across several systems or concealed within otherwise legitimate business transactions. An isolated authentication failure, an unusual transaction, or an unexpected privilege change may appear insignificant when examined independently. However, the relationship between such events can reveal a more meaningful indication of compromise. Effective SAP security therefore requires analytical methods capable of connecting individual observations, establishing normal patterns of activity, and interpreting deviations within their operational and business context.

SAP Enterprise Threat Detection provides a specialized approach to examining security-relevant information generated within SAP environments. When considered alongside Security Information and Event Management (SIEM) systems, it offers opportunities to correlate application-level evidence with security events originating from networks, operating systems, databases, and identity management infrastructure. This integration can improve the visibility available to security analysts investigating suspicious behavior across an enterprise. Nevertheless, collecting events from multiple sources does not automatically produce reliable threat detection. Differences in event formats, incomplete audit records, inconsistent timestamps, and variations in system configuration can complicate correlation. A practical detection framework must therefore address both the technical integration of security information and the interpretation of the resulting evidence.

The central research problem concerns the difficulty of distinguishing meaningful threats from legitimate operational variations across complex SAP landscapes. Rule-based monitoring can generate excessive alerts when detection conditions do not adequately reflect business activity, while narrowly defined rules may overlook unfamiliar attack patterns. Similarly, statistical deviations do not necessarily indicate malicious intent, because legitimate activities such as month-end financial processing, emergency maintenance, and authorized administrative operations can produce unusual event patterns. Treating every deviation as a threat may overwhelm investigators, whereas ignoring deviations that appear individually insignificant may allow coordinated attacks to remain undetected. This study argues that effective threat identification requires a risk-centric approach that combines event relationships, statistical evidence, and the operational significance of the affected systems.

A further research gap concerns the relationship between anomaly detection and security event prioritization. Existing monitoring practices may identify suspicious events without adequately explaining how those events contribute to broader business exposure. Detection accuracy alone cannot establish the operational value of a security monitoring system if investigators must still examine numerous unrelated alerts to identify the incidents requiring immediate attention. A more comprehensive approach should consider event severity, behavioral deviation, access privileges, asset importance, and potential business consequences. The motivation for this research is to develop a structured methodology that connects these factors, enabling SAP security teams to evaluate suspicious activity according to its evidential strength and potential impact rather than relying exclusively on isolated alert severity.

The primary objective of this study is to establish a risk-centric methodology for detecting cyber threats across SAP landscapes through security event correlation, statistical anomaly profiling, and enterprise-wide SIEM analytics. The proposed approach examines how security records from heterogeneous sources can be collected, normalized, enriched, and correlated to identify potentially significant patterns of activity. It also considers how statistical profiles of users, transactions, and system operations can support the recognition of deviations from established behavioral baselines. A further objective is to develop a consistent basis for prioritizing correlated events according to their severity, confidence, and potential business impact. These objectives are intended to support systematic investigation while preserving the distinction between suspicious behavior and confirmed security incidents.

The research is guided by several questions. How can security events originating from different SAP components and supporting infrastructure be correlated to reveal suspicious activity that isolated monitoring may overlook? To what extent can statistical anomaly profiling complement conventional rule-based detection without producing an unacceptable increase in false-positive alerts? How can event severity, behavioral deviation, and business exposure be combined to establish meaningful risk priorities? What evaluation measures can determine whether the proposed methodology improves detection precision, recall, correlation effectiveness, and incident identification time? Addressing these questions requires a systematic evaluation framework that considers both the analytical performance of detection methods and the practical demands of enterprise security operations.

The significance of this research extends to cybersecurity engineering, SAP administration, internal auditing, and enterprise risk management. By connecting application-specific security evidence with enterprise SIEM analytics, the proposed methodology seeks to improve the interpretation of security events across interconnected systems. Its potential practical contribution lies in helping analysts investigate related activities, distinguish routine operational changes from suspicious patterns, and allocate investigative resources according to risk. Its academic contribution lies in defining an evaluable approach for examining the relationship between event correlation, statistical anomaly detection, and risk-based prioritization in SAP environments. The methodology also emphasizes the importance of reproducible evaluation, transparent scoring criteria, and explicit recognition of limitations associated with incomplete data and changing operational behavior.

The study consequently establishes a foundation for evaluating risk-centric threat detection as an integrated security monitoring approach rather than a collection of independent detection rules. Its contribution is assessed through defined threat scenarios, comparative detection measures, and analysis of the operational trade-offs between sensitivity, false-positive reduction, and investigation efficiency. The proposed methodology does not assume that statistical anomalies alone can establish malicious activity or that SIEM integration eliminates existing monitoring limitations. Instead, it treats correlated evidence, behavioral context, and business risk as complementary elements in the investigation process. This perspective provides a structured basis for improving threat visibility across SAP landscapes while supporting more consistent security decisions and further research into enterprise application protection.

2. Research Background and Threat Landscape

SAP enterprise environments comprise interconnected application servers, database platforms, user interfaces, business modules, and external integration services that collectively support essential organizational activities. Their security architecture must protect business information while maintaining reliable access to authorized users and applications. Access control mechanisms, user authentication, authorization profiles, system configuration, and audit facilities form important components of this protection. Within these environments, security-relevant information is distributed across application logs,

system logs, audit trails, transaction records, and authentication histories. Each source provides a partial view of system behavior, but the combined evidence can reveal relationships that remain difficult to recognize through isolated inspection. Understanding the distribution and meaning of these records is therefore fundamental to establishing a reliable approach to SAP threat detection.

Application logs provide information about operational events, application errors, processing activities, and selected security-related conditions. Security audit records can capture configured events associated with user activity, authorization checks, changes to sensitive configurations, and other actions relevant to security investigation. Authentication records provide evidence concerning successful and unsuccessful access attempts, while system logs can help identify application failures, service interruptions, and unexpected system behavior. Transaction records offer additional context by indicating how users interact with business processes and sensitive organizational data. However, the availability and completeness of these records depend on system configuration, retention policies, logging settings, and the nature of the event being investigated. Consequently, a threat detection methodology must account for differences in coverage and evidential reliability rather than assuming that every relevant action produces a complete and consistent record.

Unauthorized access represents a significant threat because compromised credentials or inappropriate access permissions may allow individuals to enter systems or perform activities beyond their legitimate responsibilities. Privilege misuse creates additional exposure when users possess excessive authorizations, exploit legitimate access for unauthorized purposes, or modify sensitive settings without appropriate justification. Suspicious transaction activity may involve unusual financial postings, unexpected master data changes, abnormal execution sequences, or operations performed at atypical times. These activities cannot automatically be classified as malicious because legitimate business processes also generate exceptional events. Nevertheless, combinations of unusual access attempts, privilege changes, and sensitive transactions may provide stronger indications of potential compromise than any individual event considered separately. Effective analysis must therefore distinguish between isolated irregularities and patterns that collectively suggest a meaningful security concern.

Abnormal system activity introduces another dimension to the threat landscape. Repeated authentication failures, unusual transaction frequencies, unexpected communication between system components, and abrupt changes in administrative activity may indicate account misuse, unauthorized automation, or attempts to exploit weaknesses in system configuration. Certain attacks may also involve multiple stages, with initial access followed by privilege escalation, reconnaissance, data modification, or unauthorized extraction. Because these stages can occur across different components, their evidence may be dispersed among several logging facilities. Traditional threshold-based monitoring can identify recognizable symptoms, but it may fail to connect events when individual actions remain below predefined thresholds. Statistical anomaly profiling can complement these techniques by examining deviations in event frequency, timing, sequence, and user behavior, provided that the resulting indicators are interpreted within an appropriate operational context.

Security Information and Event Management systems provide a centralized means of collecting, normalizing, storing, and analyzing security events from heterogeneous sources. Their capabilities commonly include rule-based correlation, alert generation, event filtering, reporting, and support for incident investigation. Correlation rules can associate events through shared attributes such as user identifiers, source addresses, timestamps, affected systems, or event categories. These relationships help analysts reconstruct activity that would otherwise remain fragmented across separate records. However, the usefulness of centralized monitoring depends on the quality of collected data, the appropriateness of correlation conditions, and the availability of contextual information. In SAP environments, general infrastructure alerts may not explain the business significance of a particular transaction, authorization

change, or application-level operation. Incorporating SAP-specific event meaning into enterprise monitoring is therefore essential to interpreting correlated evidence accurately.

Anomaly detection techniques provide a complementary means of identifying activity that differs from established patterns. Statistical approaches can examine event counts, transaction distributions, access frequency, timing variations, and deviations from user or system baselines. Simple statistical thresholds may be useful where activity follows relatively stable patterns, while more elaborate profiling can accommodate differences between users, roles, transaction categories, and operating periods. Yet anomalous behavior is not equivalent to malicious behavior. Seasonal processing, organizational changes, emergency maintenance, and legitimate administrative work may produce substantial deviations from normal conditions. Detection methods must consequently consider baseline quality, the frequency of legitimate exceptions, and the possibility that malicious activity may gradually resemble ordinary behavior. Combining statistical indicators with event correlation and rule-based knowledge offers a more balanced basis for investigation than depending exclusively on any single detection method.

A key research gap lies in establishing a consistent relationship between SAP-specific security evidence and the broader analytical capabilities of enterprise SIEM platforms. Centralized collection improves visibility, but it does not by itself ensure that correlated events receive appropriate interpretation or risk prioritization. Generic event rules may lack the business context needed to distinguish an authorized exceptional transaction from a potentially fraudulent operation. Conversely, SAP-specific monitoring may provide detailed application evidence without sufficient visibility into related infrastructure activity, identity events, or network behavior. Differences in event schemas, clock synchronization, identity representation, and retention practices further complicate cross-system correlation. A useful methodology must address these limitations through event normalization, contextual enrichment, correlation across relevant sources, and risk assessment that reflects both technical severity and potential business consequences.

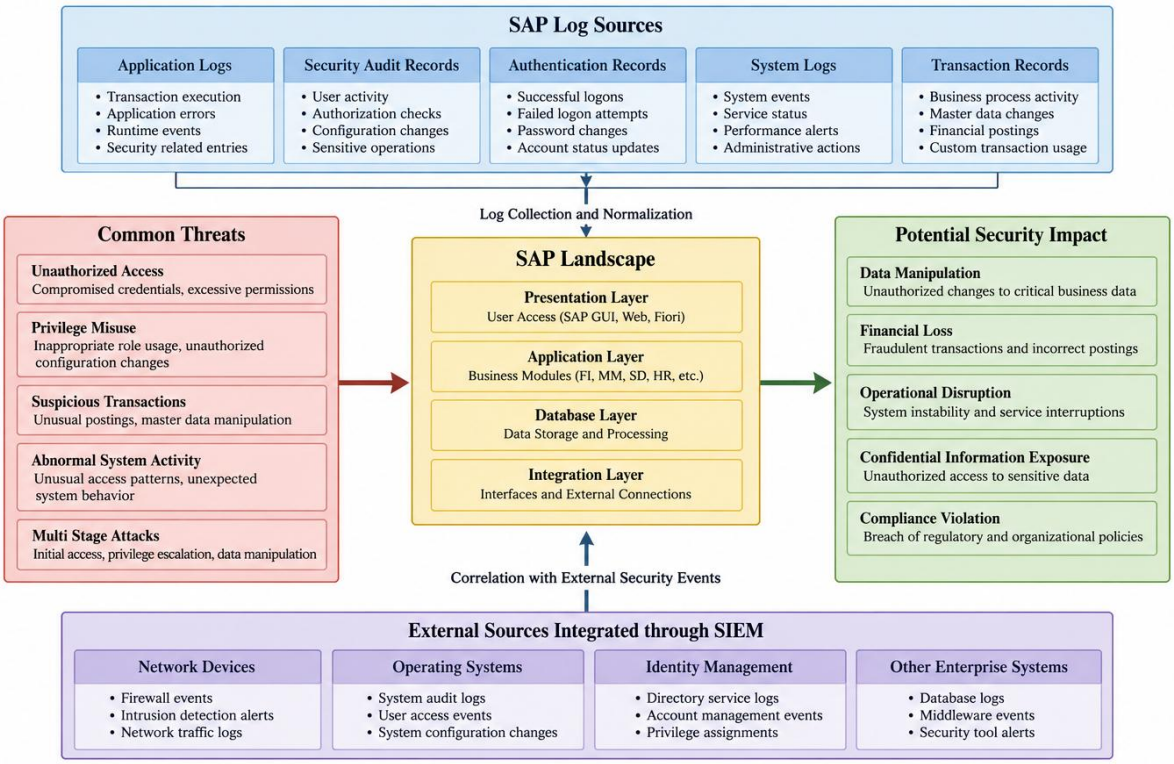


Figure 1: SAP Enterprise Security Threat Landscape and Log Sources

Effective threat identification across interconnected SAP systems therefore requires several complementary capabilities. First, logging must be configured to capture relevant security events, with appropriate controls over retention, access, and integrity. Second, collected records must be normalized without discarding meaningful application-specific attributes. Third, correlation mechanisms must connect events using reliable identifiers and reasonable temporal relationships while accounting for missing or delayed records. Fourth, statistical profiles must reflect legitimate differences in user responsibilities, transaction patterns, and business schedules. Fifth, risk assessment must distinguish detection confidence from potential impact, ensuring that high-consequence events receive appropriate investigative attention without treating every anomaly as a confirmed incident. Finally, evaluation must consider detection precision, recall, false-positive rates, identification time, and investigative usefulness. These requirements establish the research foundation for the risk-centric methodology developed in the following section.

3. Proposed Risk-Centric Detection Methodology

The proposed risk-centric detection methodology establishes a structured approach for identifying, correlating, and prioritizing potential cyber threats across interconnected SAP landscapes. Its central premise is that reliable threat detection depends not only on recognizing individual suspicious events but also on understanding their relationships, behavioral context, and potential consequences for business operations. The methodology combines SAP security event collection, data normalization, contextual enrichment, statistical anomaly profiling, risk assessment, and investigation support within a unified analytical process. SAP Enterprise Threat Detection capabilities provide an important source of application-specific security evidence, while enterprise Security Information and Event Management (SIEM) systems support the integration of relevant events from infrastructure and external security controls. The resulting framework is intended to improve the interpretation of fragmented security information without assuming that every irregular event represents a confirmed attack.

The research framework consists of six connected stages: event acquisition, normalization, contextual enrichment, event correlation, anomaly assessment, and risk-based investigation. Event acquisition gathers relevant records from SAP application servers, security audit facilities, authentication services, database platforms, network devices, operating systems, and other supporting systems. Normalization converts these records into a consistent analytical representation while retaining source-specific attributes needed for subsequent investigation. Contextual enrichment associates events with available information about user identities, assigned roles, affected assets, business transactions, and system importance. Correlation procedures then examine relationships among events, while anomaly assessment identifies deviations from established patterns of legitimate activity. Finally, risk-based investigation combines the available evidence to determine which alerts warrant immediate attention, further analysis, or routine observation.

Event acquisition requires careful consideration of the security information available from each component. Application logs can reveal processing failures and unusual application behavior, while security audit records may provide evidence of configured security-relevant actions, authorization checks, and changes to sensitive settings. Authentication records contribute information about successful access, failed login attempts, and account-related activity. Transaction records can identify unusual operations involving financial postings, master data, or other business processes. Infrastructure logs provide complementary evidence about network connections, operating system activity, and supporting services.

Collection procedures should account for logging configuration, record retention, source reliability, and differences in event availability. Missing records must not be interpreted automatically as evidence of malicious activity, because incomplete logging, system interruptions, and configuration limitations can produce similar observations.

Normalization is necessary because heterogeneous sources frequently represent equivalent information in different formats. Event timestamps, user identifiers, system names, severity levels, transaction references, and source addresses may differ across platforms. The proposed methodology maps these attributes into a common event structure while preserving the original record and its provenance. Timestamp normalization supports chronological reconstruction, provided that the original timestamp, source time zone, and any known synchronization limitations remain available for verification. Identity resolution associates records with the appropriate user or technical account without incorrectly merging distinct identities. Event classification assigns records to relevant categories, such as authentication, authorization, configuration change, transaction execution, and infrastructure activity. These procedures create a consistent foundation for correlation while reducing the risk of losing details that may become important during forensic examination.

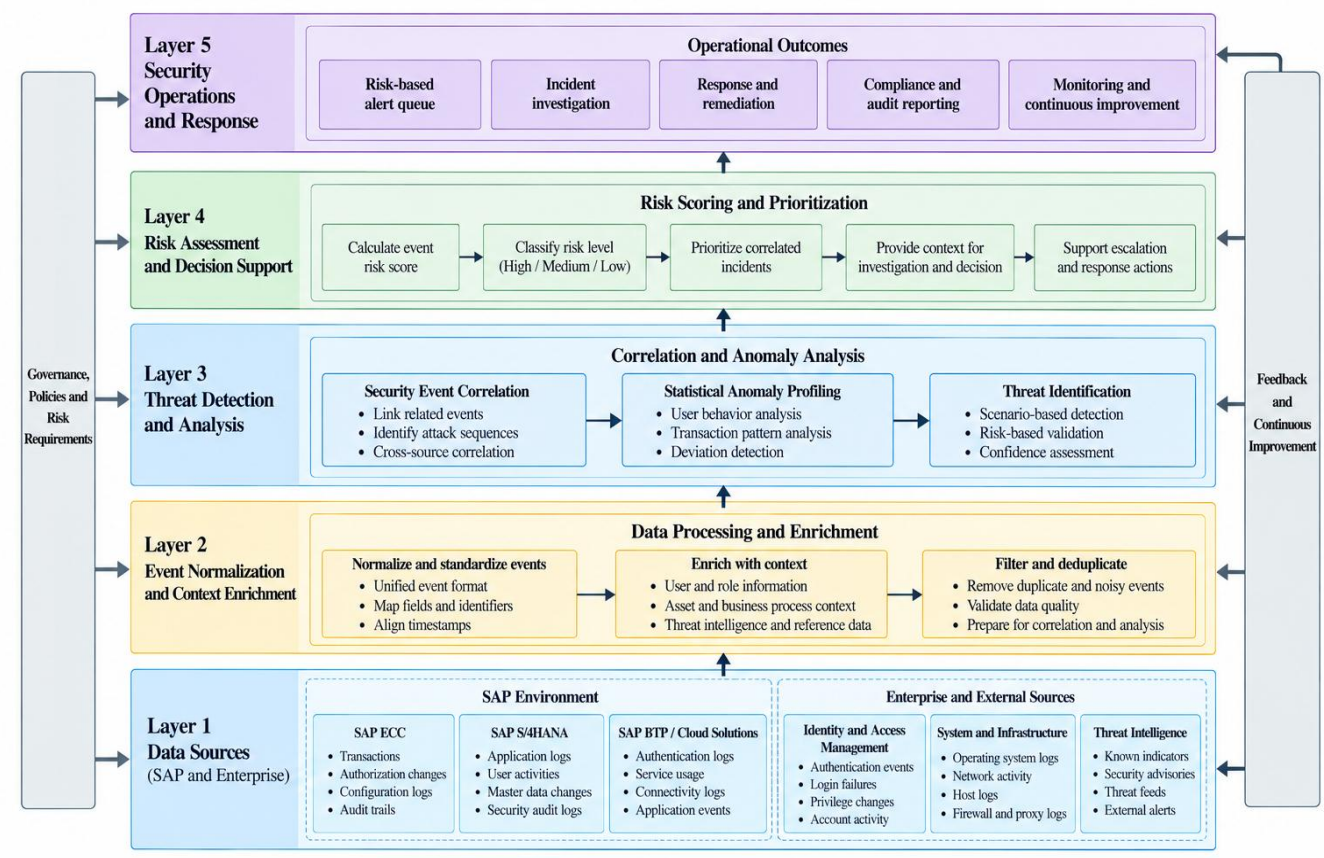


Figure 2: Layered Architecture of the Proposed Risk-Centric Detection Framework

Contextual enrichment adds information required to interpret an event beyond its technical description. A failed authentication attempt, for example, may have limited significance when considered alone, but its importance may increase when followed by successful access, unusual privilege changes, and activity involving sensitive transactions. Similarly, an exceptional transaction may be legitimate when performed

by an authorized employee during a scheduled business process, yet warrant investigation when associated with an unfamiliar access pattern or an unexplained modification to user permissions. The methodology therefore considers user responsibilities, authorization assignments, transaction sensitivity, asset importance, historical activity, and relevant operational schedules where such information is available. Enrichment must remain evidence-based: missing contextual information should be recorded as uncertainty rather than replaced with unsupported assumptions about user intent or business legitimacy.

Security event correlation examines whether separate records form a meaningful sequence or share attributes that indicate a possible relationship. The proposed approach uses temporal proximity, common user or account identifiers, affected system references, source addresses, transaction relationships, and event categories to identify related activity. Rule-based correlation can detect known sequences, including repeated access failures followed by successful authentication and subsequent privileged operations. Statistical correlation can supplement these rules by identifying unusual combinations of event frequency, timing, or sequence that fall outside established operating patterns. Correlation windows must be selected according to the nature of the activity, because overly narrow intervals may separate related events, while excessively broad intervals may associate unrelated operations. Where records lack reliable identifiers or timestamps, the resulting relationship should receive lower confidence. The objective is to reconstruct plausible activity sequences while maintaining a clear distinction between observed facts and analytical inference.

Statistical anomaly profiling establishes reference patterns against which current activity can be assessed. Profiles may be constructed for individual users, technical accounts, transaction categories, administrative operations, and system components. Depending on the available data, useful indicators include event frequency, transaction volume, access time, failed authentication counts, changes in privilege usage, and deviations from historical distributions. Basic statistical measures, including means, standard deviations, percentiles, and frequency thresholds, can identify departures from established behavior without requiring an unnecessarily complex analytical model. Separate baselines may be appropriate for different user roles or business periods because a single enterprise-wide baseline can obscure legitimate differences in operational activity. Baselines should also be reviewed when organizational responsibilities, application configurations, or business schedules change. An anomaly indicates that activity merits examination, not that malicious behavior has been conclusively established.

Risk assessment combines the severity of the observed event, the magnitude of behavioral deviation, the importance of the affected asset, and the potential business impact into a structured prioritization measure. For methodological consistency, these factors should be converted to documented scales before they are combined. One possible formulation is:

$$R = w_s S + w_a A + w_i I + w_b B$$

where R denotes the composite risk score, S represents event severity, A represents anomaly magnitude, I represents asset importance, and B represents potential business impact. The weights w_s, w_a, w_i, and w_b express the relative contribution of each factor and should be selected and validated against the intended operational environment. If each input is normalized to a common scale, the resulting score can support consistent comparison among alerts. However, the score is a prioritization aid rather than a calibrated probability of attack. Critical indicators, such as evidence of unauthorized privileged access to a sensitive system, may also justify mandatory escalation regardless of the composite score. This arrangement reduces the risk that a serious event will be obscured by comparatively low values assigned to other factors.

The final stage translates correlated evidence and risk scores into an investigation workflow. High-priority alerts should provide analysts with the underlying records, event sequence, affected systems, relevant identity and authorization context, contributing anomaly indicators, and reasons for the assigned priority.

Medium-priority alerts may require additional evidence gathering, comparison with historical activity, or confirmation from system owners, while lower-priority observations can be retained for trend analysis and subsequent correlation. Analysts should be able to trace every alert to its original records and understand which rules, statistical indicators, and contextual attributes influenced its classification. Investigation outcomes can then support refinement of correlation rules, correction of inaccurate asset classifications, and adjustment of behavioral baselines. This feedback process must be governed carefully to prevent incorrect analyst decisions from becoming permanent assumptions within the detection system. Overall, the methodology provides a repeatable basis for examining how integrated security evidence and risk-sensitive prioritization can improve threat identification across SAP landscapes. Its effectiveness must ultimately be established through controlled evaluation using documented threat scenarios, reliable reference labels, and measurable detection and investigation outcomes.

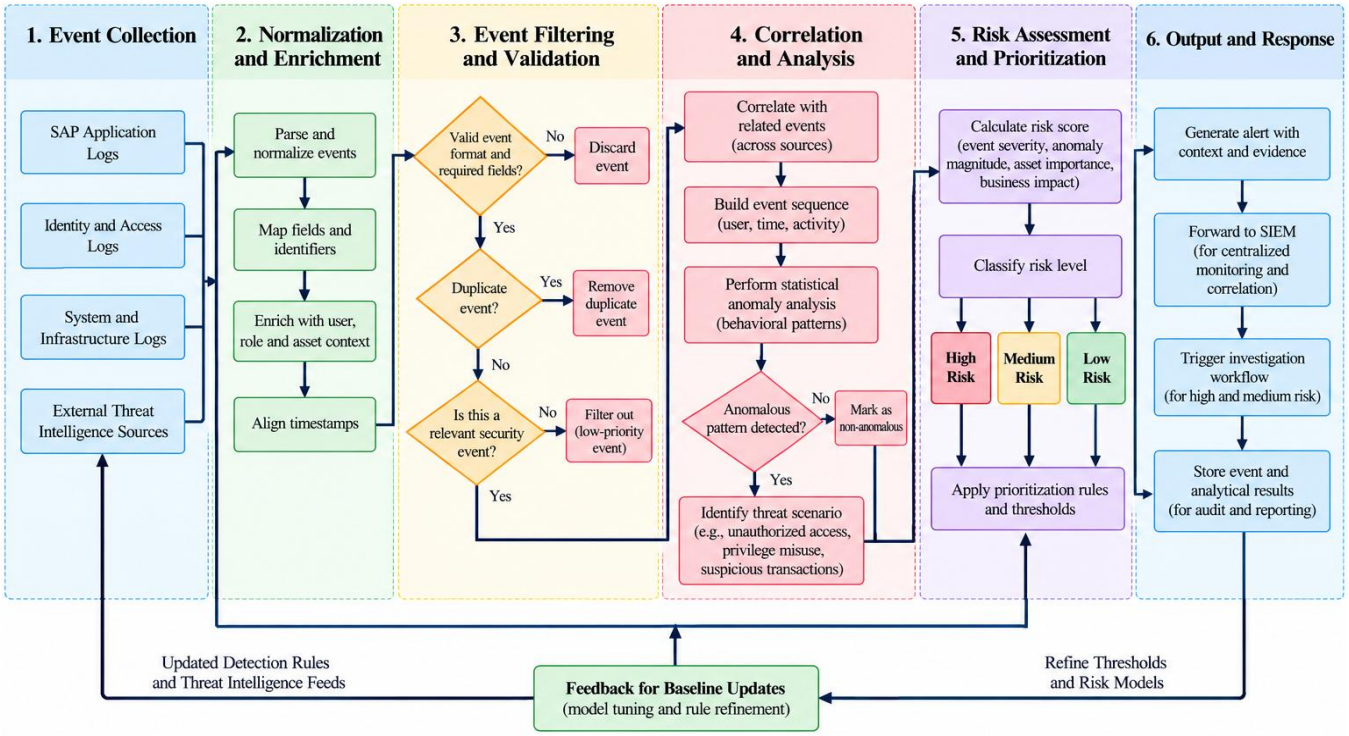


Figure 3: Security Event Processing and Correlation Workflow

4. Statistical Anomaly Profiling and SIEM Integration

Statistical anomaly profiling provides a systematic means of identifying security-relevant activities that differ from established patterns of behavior across SAP landscapes. Its value lies in extending conventional rule-based monitoring beyond predefined conditions to examine variations in user activity, transaction frequency, access timing, and system operations. Within the proposed methodology, anomaly profiling complements SAP Enterprise Threat Detection and enterprise Security Information and Event Management (SIEM) capabilities by supplying behavioral evidence that can be evaluated alongside audit records and infrastructure events. The objective is not to classify every deviation as a cyber threat, but to determine whether an observed variation warrants further investigation when considered with its surrounding context. This distinction is particularly important in enterprise applications, where legitimate operational exceptions can resemble suspicious behavior.

Behavioral baselines form the foundation of this process. A baseline represents an established pattern of activity against which subsequent observations can be compared. User-level profiles may describe typical login periods, transaction categories, authentication frequency, and the extent of privileged operations. Transaction-level profiles can characterize normal posting volumes, transaction sequences, and patterns associated with particular business functions. System-level profiles may capture application activity, administrative operations, service behavior, and recurring event frequencies. These profiles should reflect meaningful differences among users, technical accounts, business roles, and operating schedules rather than treating all activity as statistically equivalent. For example, an administrator performing a scheduled maintenance task should be evaluated against an appropriate administrative baseline rather than the profile of an ordinary business user. Baseline construction must therefore combine statistical evidence with an understanding of legitimate responsibilities and operational conditions.

Several statistical techniques can support the identification of deviations from these reference patterns. Frequency distributions and rolling averages can reveal unusual increases in authentication failures, privileged actions, or transaction volumes. Standard deviations can identify observations that depart substantially from a relatively stable distribution, while percentiles can provide useful thresholds when activity is unevenly distributed or affected by extreme values. Median-based measures may offer greater resistance to unusually large observations than arithmetic averages. Time-based comparisons can identify activities that occur outside customary operating periods, while sequence analysis can reveal unusual combinations of otherwise familiar events. The choice of method should depend on the available data, baseline stability, and the characteristics of the activity being examined. Where event counts are sparse or distributions change considerably over time, simple thresholds may produce unreliable conclusions and should be supplemented with contextual checks.

A practical implementation should distinguish ordinary variation from deviations that merit security investigation. Consider a user whose transaction volume increases substantially during a financial closing period. A general enterprise-wide threshold may classify this increase as anomalous, even though the activity is consistent with established business requirements. By contrast, a smaller increase in privileged operations accompanied by an unfamiliar authentication pattern and an unauthorized configuration change may carry greater security significance. This comparison illustrates why statistical magnitude alone is insufficient for reliable detection. The methodology evaluates anomalies in relation to user responsibilities, transaction sensitivity, asset importance, and related security events. Baselines should also be periodically reviewed to account for changes in organizational structure, transaction patterns, application configuration, and scheduled business activity. Such reviews help prevent outdated behavioral expectations from generating persistent false-positive alerts or overlooking gradually changing activity.

SIEM integration extends the analysis by connecting SAP-specific records with information from supporting infrastructure and security controls. SAP audit events may describe user actions or authorization-related changes, whereas operating system logs, network records, identity management events, and database logs provide complementary evidence about the surrounding technical environment. Correlation across these sources can help reconstruct sequences that are difficult to identify within a single application. For example, an unusual authentication event may become more significant when followed by a privilege modification, an unexpected connection from a network segment, and access to a sensitive business function. The strength of this inference depends on whether the records can be reliably associated with the same account, system, or activity sequence. Correlation procedures must therefore account for the quality of identity mappings, event timestamps, source attributes, and the possibility that several unrelated activities share common characteristics.

The proposed approach combines rule-based detection with statistical anomaly indicators to balance known threat recognition and the identification of less familiar patterns. Rules remain useful for recognizable conditions, such as repeated failed logins, unauthorized changes to sensitive configurations,

or suspicious sequences of privileged activity. Statistical profiling supplements these rules by highlighting deviations that do not necessarily match a predefined signature. When both methods identify related evidence, the combined alert can receive greater investigative attention, provided that the underlying observations are sufficiently independent and reliable. However, two detection methods reporting the same underlying event do not automatically constitute two independent pieces of evidence. The framework should therefore preserve the provenance of each indicator, document the reason for correlation, and avoid inflating confidence simply because multiple rules respond to a common source record. This approach supports more transparent and defensible security decisions.

Risk scoring, alert deduplication, and false-positive reduction determine how analytical findings are translated into manageable investigation priorities. Each correlated event group can be assessed using event severity, anomaly magnitude, affected asset importance, and potential business impact. These factors should be normalized to documented scales and combined using weights appropriate to the intended operational environment. The resulting score supports comparative prioritization but should not be interpreted as a calibrated probability of compromise unless it has been validated for that purpose. Deduplication reduces repeated notifications generated by recurring records, overlapping detection rules, or multiple monitoring components reporting the same underlying activity. Effective deduplication must preserve important differences between events, especially when repeated activity indicates persistence, escalation, or an evolving attack. False-positive reduction should rely on improved context, better baselines, and refined correlation logic rather than simply suppressing alerts that create a high investigative workload.

Reliable integration also depends on data quality and the consistency of event processing. Heterogeneous sources may use different field names, severity classifications, timestamp formats, time zones, and identity conventions. Normalization must reconcile these differences without discarding source-specific details required for verification. Time synchronization is particularly important because event ordering often determines whether separate observations can reasonably be interpreted as part of the same activity sequence. Where clocks are inconsistent or records arrive late, the system should preserve original timestamps and record any adjustments applied during processing. Log completeness requires similar attention: disabled auditing, expired retention periods, interrupted collection, and unavailable source systems can leave gaps in the evidence. Such gaps should be represented explicitly, since the absence of a record cannot establish that an action did not occur. Data quality indicators can help analysts interpret correlation results in light of these limitations.

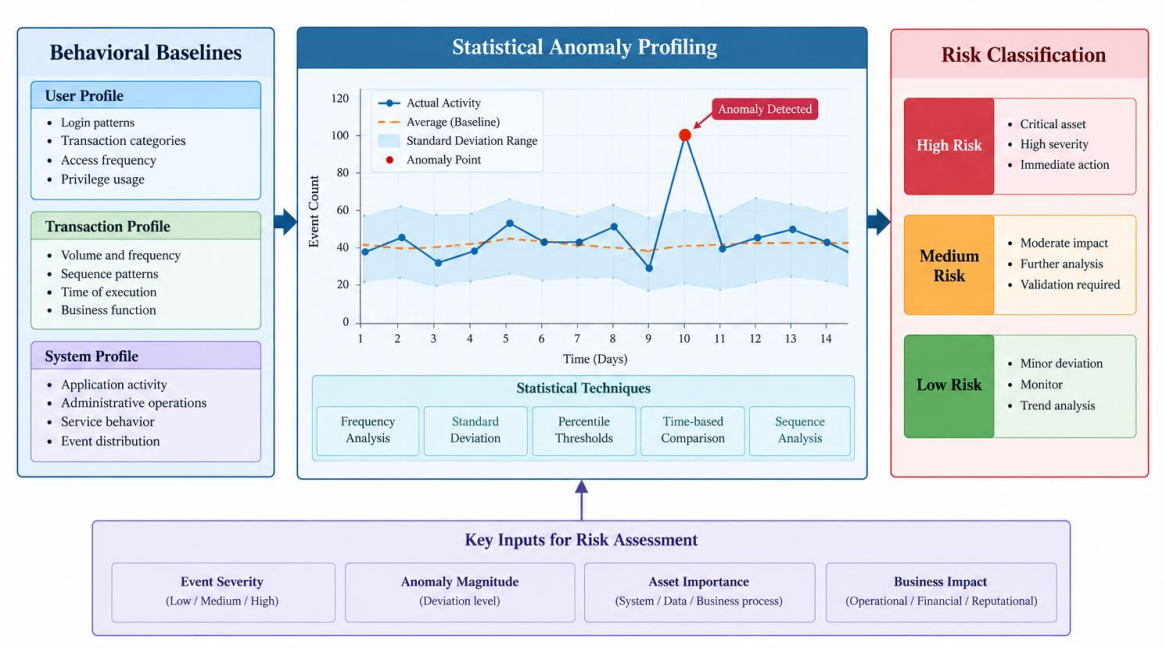


Figure 4: Statistical Anomaly Profiling and Risk Classification Model

The effectiveness of statistical anomaly profiling and SIEM integration must ultimately be evaluated through measurable outcomes rather than assumed from architectural design alone. Relevant measures include detection precision, recall, false-positive rate, detection latency, correlation accuracy, and the time required to investigate an alert. Comparative testing should examine whether the combined methodology identifies suspicious sequences that conventional rules miss, while also assessing whether the added analytical complexity produces excessive false alarms. Results may vary according to event quality, baseline selection, system configuration, threat scenario, and the availability of reliable reference labels. For this reason, evaluation should include both routine operational variations and deliberately constructed suspicious scenarios, with the assumptions and limitations documented clearly. The proposed integration offers a structured basis for improving SAP threat monitoring, but its practical value depends on disciplined configuration, transparent analytical decisions, and validation within the operational conditions of the enterprise.

5. Experimental Design and Evaluation Criteria

The experimental design establishes a systematic basis for evaluating the proposed risk-centric methodology under security conditions representative of SAP enterprise environments. The primary purpose is to determine whether combining security event correlation, statistical anomaly profiling, and risk-based prioritization improves threat identification compared with conventional rule-based monitoring. The evaluation must distinguish between the conceptual advantages of the proposed architecture and performance improvements demonstrated through measurable evidence. Accordingly, the research design emphasizes controlled test scenarios, documented event characteristics, consistent detection criteria, and repeatable assessment procedures. The evaluation considers both technical outcomes, such as detection precision and identification latency, and operational outcomes, including alert quality and the usefulness of correlated evidence for incident investigation. No performance advantage is assumed in advance; the proposed methodology must demonstrate its value through comparative testing.

A controlled comparative design is appropriate because it allows two detection approaches to be assessed against the same underlying event sequences. The first approach represents conventional rule-based monitoring, in which predefined conditions generate alerts when specified event patterns or thresholds are encountered. The second approach implements the proposed risk-centric methodology, combining rule-based indicators with statistical deviations, contextual enrichment, cross-source correlation, and risk scoring. Both approaches should receive equivalent source events and operate under documented configuration settings to avoid attributing differences in input quality to the detection method itself. Where separate configuration requirements cannot be eliminated, their influence should be recorded explicitly. The experimental design should also preserve a distinction between event detection and incident classification, since identifying an unusual record does not necessarily establish that the record belongs to a genuine security incident.

The test scenarios should cover several representative threat categories. Unauthorized access scenarios may include repeated failed authentication attempts followed by successful access from an unfamiliar source, or account activity inconsistent with established access patterns. Privilege misuse scenarios may involve unexpected authorization changes, sensitive administrative operations, or privileged activity outside an established maintenance window. Unusual transaction scenarios may include atypical posting volumes, unexpected changes to critical master data, and transactions inconsistent with the normal responsibilities of the associated user. Suspicious event-sequence scenarios should examine combinations of individually inconclusive observations, such as unusual authentication, subsequent privilege modification, and access to a sensitive business function. Benign comparison scenarios are equally important and should include legitimate month-end processing, authorized emergency maintenance, routine administrative activity, and temporary increases in transaction volume. Including both suspicious and legitimate exceptions helps determine whether the methodology can distinguish meaningful security indicators from normal operational variation.

Data preparation requires a consistent representation of events before either detection approach is evaluated. Records from SAP application logs, security audit facilities, authentication services, system logs, and relevant infrastructure sources should be mapped to a common event structure containing available timestamps, source identifiers, user or account references, event categories, affected systems, and relevant transaction attributes. Original records should be retained so that normalization decisions can be traced and the analytical process independently reviewed. Duplicate records, missing fields, inconsistent timestamps, and incomplete event sequences must be identified and handled according to documented rules. Missing information should not automatically be interpreted as evidence of malicious activity. Event labels should be assigned using predefined scenario descriptions and, where feasible, independent review of the underlying records. The distinction between a suspicious event, a confirmed incident, and an ordinary operational exception must remain explicit throughout the evaluation.

Baseline construction should use a defined reference period that represents legitimate activity and is separated, where practicable, from the observations used to evaluate detection performance. User profiles can capture authentication frequency, customary access periods, transaction categories, and privileged operations. Transaction profiles can describe typical volumes and execution patterns, while system profiles can characterize recurring administrative events and application activity. Baselines should account for relevant differences among user roles, business schedules, and system functions. Where data are divided into training, validation, and test subsets, the division should prevent information from the evaluation period from influencing the construction of the reference profiles. This separation reduces the risk of data leakage and provides a more credible assessment of detection behavior on previously unseen observations. The evaluation assumptions should specify the baseline duration, event inclusion criteria, threshold selection procedure, and treatment of legitimate operational changes.

The comparative evaluation should examine conventional rule-based monitoring and the proposed risk-centric approach under equivalent test conditions. The rule-based method should use documented detection rules and thresholds representative of the monitoring configuration being evaluated. The proposed method should apply the same underlying rules where appropriate, supplemented by statistical anomaly indicators, contextual information, event correlation, and risk-based alert prioritization. This arrangement helps isolate the contribution of the additional analytical components without implying that every rule-based system has identical capabilities. Performance should be reported separately for individual threat scenarios and for the complete test set. Where possible, an additional ablation analysis should examine the contribution of statistical profiling, contextual enrichment, and cross-source correlation by evaluating the framework with selected components disabled. Such analysis can clarify whether observed improvements arise from a particular component or from their combined operation.

Detection performance should be assessed through established quantitative measures. Precision represents the proportion of generated positive alerts that correspond to genuinely positive cases, while recall represents the proportion of positive cases successfully identified. The false-positive rate measures the proportion of benign cases incorrectly classified as positive. These measures provide complementary information because a detection method can increase recall by generating more alerts while simultaneously increasing the number of false positives. Detection latency should measure the elapsed time between the earliest observable evidence of a defined test incident and the generation of an actionable alert, with the starting point specified consistently across methods. Event correlation effectiveness should assess whether related records are grouped into the correct activity sequences, using scenario-level reference labels and documented matching criteria. Investigation efficiency may additionally be examined through the number of duplicate alerts, the time required to reconstruct an incident, and the completeness of supporting evidence presented to analysts.

Reproducibility and validity require the experimental procedures to be documented in sufficient detail for another researcher to repeat the assessment. The research record should describe the event sources, preprocessing rules, scenario definitions, labeling criteria, baseline construction, detection thresholds, risk-scoring weights, correlation windows, and software configuration used in the evaluation. If real enterprise records are unavailable, controlled synthetic or sanitized event data may provide a practical alternative, provided that their construction and limitations are disclosed. Synthetic scenarios can support repeatable testing but may not fully represent the irregularities, incomplete logging, and changing behavior encountered in production systems. Internal validity may also be affected by inconsistent configurations or data leakage, while external validity depends on whether the tested scenarios adequately represent different SAP configurations and operational conditions. These limitations should be acknowledged when interpreting the results rather than treated as minor implementation details.

The evaluation framework therefore provides a structured means of determining whether risk-centric analysis offers measurable benefits beyond conventional monitoring. Its central comparison concerns the ability to identify suspicious activity, associate related events, distinguish legitimate deviations from potential threats, and prioritize alerts according to their evidential and operational significance. Results should include the underlying counts of true positives, false positives, true negatives, and false negatives alongside derived performance metrics, allowing readers to assess how the reported values were obtained. Where repeated trials are conducted, variation across runs should also be reported, and statistical comparisons should be used when the experimental design supports them. No measured findings should be inferred solely from the proposed architecture or from illustrative scenarios. By maintaining consistent inputs, transparent assumptions, and clearly defined evaluation criteria, this design establishes a defensible foundation for the comparative performance analysis presented in the following section.

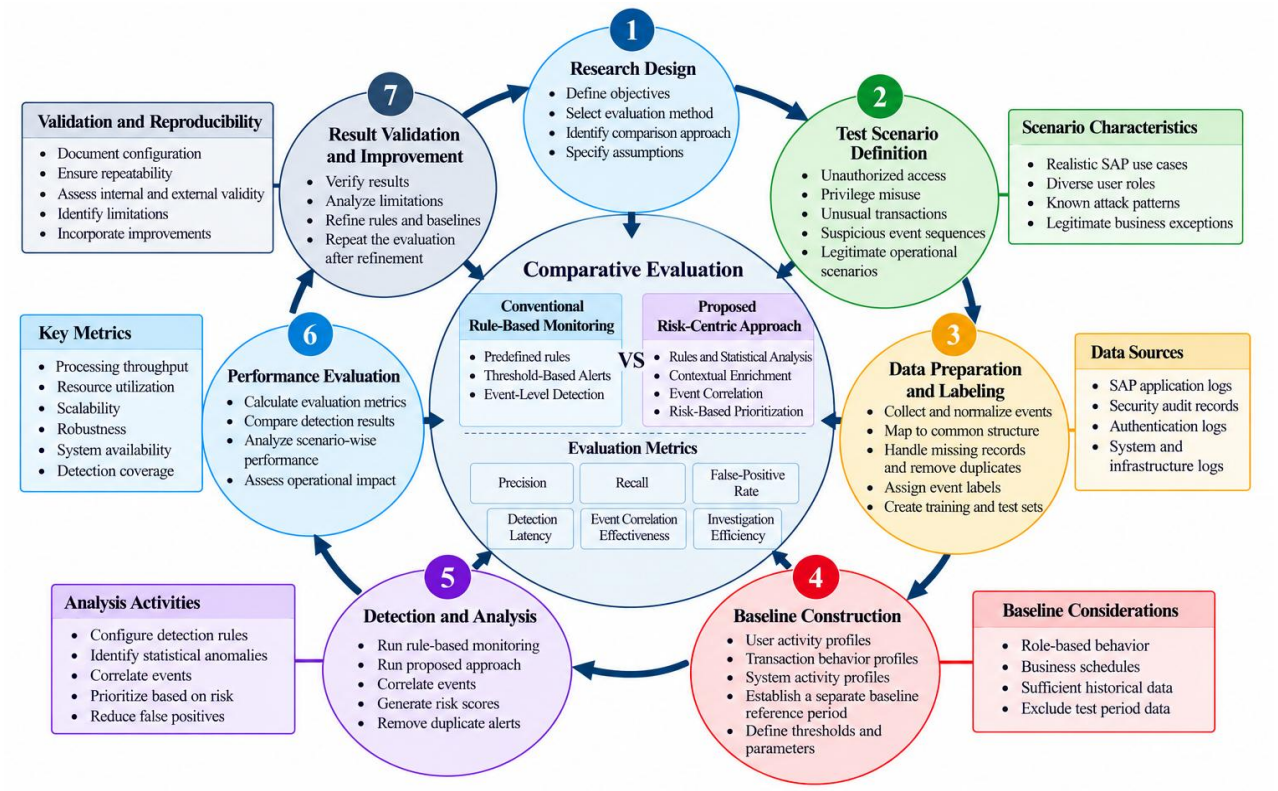


Figure 5: Experimental Evaluation Workflow and Detection Assessment Process

6. Results and Comparative Performance Analysis

The evaluation of the proposed risk-centric methodology is intended to establish whether combining security event correlation, statistical anomaly profiling, and risk-based prioritization provides practical advantages over conventional rule-based monitoring across SAP landscapes. The analysis focuses on detection precision, recall, false-positive rate, detection latency, correlation effectiveness, and investigation efficiency. These measures represent different aspects of detection performance and must be interpreted together rather than treated as interchangeable indicators of success. Conventional rule-based monitoring provides the reference approach, while the proposed methodology adds behavioral baselines, contextual enrichment, and cross-source event relationships to the detection process. Since no experimental dataset or measured performance results have been supplied, numerical improvements cannot be established at this stage. The following analysis therefore defines how the comparative findings should be presented and interpreted once the experiments described in Section 5 have been completed.

The first comparison concerns detection performance across the defined threat scenarios, including unauthorized access, privilege misuse, unusual transaction activity, and suspicious event sequences. Conventional rules may perform effectively when observed activity matches predefined conditions, such as repeated authentication failures or unauthorized configuration changes. Their performance may be less consistent when a threat is distributed across multiple event sources or involves individually inconclusive actions. The risk-centric approach is designed to examine these relationships by combining event characteristics with behavioral deviations and contextual evidence. Its comparative advantage should be established by measuring the proportion of correctly identified threat cases under identical input conditions. Results should be reported separately for each scenario because an aggregate detection score may conceal weaknesses affecting particular threat categories. Any observed improvement must be supported by the corresponding event counts and evaluation metrics.

Detection precision and recall provide complementary measures for interpreting the results. Precision indicates how many generated positive alerts correspond to genuinely positive cases, whereas recall measures how many positive cases are successfully identified. A system that generates numerous alerts may achieve high recall while placing a considerable investigative burden on security analysts. Conversely, a conservative detection configuration may reduce false alarms but overlook important incidents. The proposed methodology seeks to balance these competing objectives by evaluating correlated evidence and contextual risk before assigning investigation priority. The comparative analysis should therefore examine whether changes in recall are accompanied by changes in precision and whether the overall balance improves across the test scenarios. The final report should include the true-positive, false-positive, true-negative, and false-negative counts used to calculate each measure, allowing readers to verify the reported findings.

False positives and missed detections require particular attention because they influence both security effectiveness and operational workload. False positives can arise when legitimate business activity deviates from an established baseline, such as during financial closing, emergency maintenance, or authorized administrative operations. The proposed methodology may help distinguish these cases by incorporating user roles, transaction context, system importance, and related events. However, contextual enrichment cannot eliminate ambiguity when source records are incomplete or legitimate activity is poorly documented. Missed detections may occur when malicious actions remain within established behavioral patterns, when relevant records are unavailable, or when correlation rules fail to associate related events. The evaluation should consequently examine false-positive and false-negative cases individually, identify the conditions responsible for each error, and determine whether improvements require changes to baselines, event collection, correlation logic, or risk thresholds.

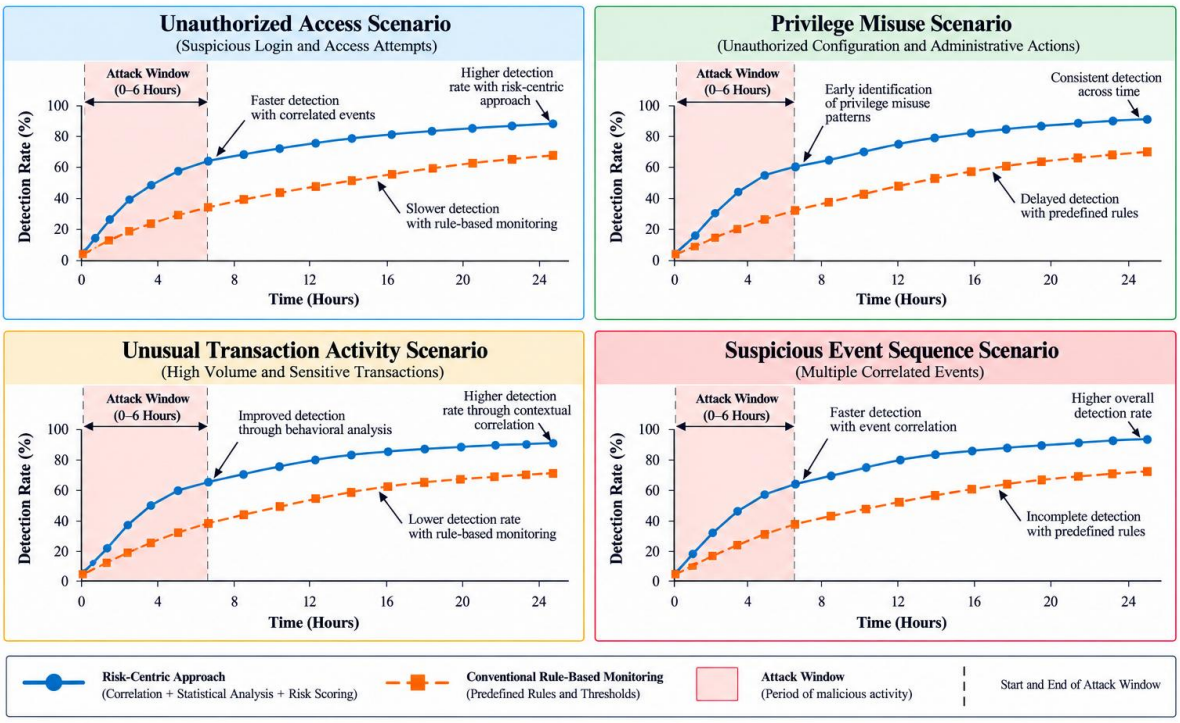


Figure 6: Comparative Detection Performance Across Threat Scenarios

Detection latency represents another important dimension of comparative performance. In enterprise security operations, identifying suspicious activity promptly can reduce the time available for an attacker to misuse compromised credentials, expand privileges, or modify sensitive business information. Latency should be measured consistently from a clearly defined starting point, such as the earliest observable event associated with a test incident, to the generation of an actionable alert. The analysis should distinguish processing time from delays caused by event collection, log delivery, normalization, and correlation windows. The proposed methodology may require additional processing to evaluate statistical deviations and enrich event context, potentially increasing computational overhead. Conversely, grouping related events into a coherent incident may reduce the time analysts spend reconstructing activity from separate alerts. The final comparison should establish whether any additional processing cost is offset by faster identification or more efficient investigation, rather than assuming that greater analytical complexity automatically improves response time.

Event correlation effectiveness and alert prioritization should be assessed in relation to the quality of the evidence available to investigators. Effective correlation connects related records into plausible activity sequences while preserving the source information needed for verification. Evaluation should determine whether the method correctly associates events belonging to the same test incident, avoids combining unrelated activities, and retains important indicators of escalation or persistence. Alert prioritization should then be assessed by examining whether high-risk scenarios receive appropriate attention relative to lower-consequence observations. The composite risk score can support this process by considering event severity, anomaly magnitude, asset importance, and potential business impact. Nevertheless, a high score does not establish malicious intent, and a low score does not guarantee that an event is harmless. The operational value of prioritization should be demonstrated through measurable reductions in duplicate alerts, investigation effort, or time required to identify significant incidents.

The interpretation of comparative results must also account for methodological limitations and operational trade-offs. Performance may vary according to the quality of SAP audit records, the availability of infrastructure logs, the duration of the behavioral reference period, and the representativeness of the test scenarios. Synthetic event sequences can support controlled and repeatable comparisons, but they may not capture all the irregularities found in production environments. Statistical baselines may also become less reliable when business practices change or when legitimate activity is sparsely represented in the reference data. Similarly, aggressive alert deduplication may reduce notification volume while inadvertently concealing repeated actions that indicate an evolving attack. These considerations make it necessary to report not only aggregate metrics but also scenario-specific errors, processing overhead, and data limitations. Where repeated trials are available, variation across runs should be examined to determine whether the observed differences are consistent rather than attributable to chance or particular test conditions.

The implications for SAP security operations depend on the measured balance between detection quality, analytical overhead, and investigative usefulness. A successful implementation would allow security analysts to examine application events alongside authentication, infrastructure, and system records, providing a broader basis for understanding suspicious activity. Risk-based prioritization could help allocate investigative resources toward events affecting sensitive systems or critical business processes, while contextual evidence could reduce the need to examine numerous disconnected alerts individually. These benefits would depend on appropriate audit configuration, reliable identity mapping, accurate asset classification, and continuous review of detection rules and behavioral baselines. Integration with enterprise SIEM platforms should therefore be evaluated as an operational capability rather than simply a technical connection between products. The results should establish whether the additional analytical components provide sufficient measurable value to justify their configuration, maintenance, and processing requirements.

The final assessment should consolidate the comparative findings into an evidence-based judgment about the proposed methodology. The results table should present the measured performance of conventional rule-based monitoring and the risk-centric approach using identical evaluation definitions, while the accompanying chart should illustrate the principal differences across the tested threat scenarios. Precision, recall, false-positive rate, detection latency, correlation effectiveness, and investigation efficiency should be reported only where the underlying observations support their calculation. Any missing metric should be identified explicitly rather than estimated without evidence. The analysis should also distinguish improvements demonstrated within the controlled evaluation from benefits that remain to be verified in production SAP environments. This approach provides a defensible basis for determining whether risk-centric correlation improves enterprise threat detection and identifies the conditions under which the methodology is most useful. It also establishes the empirical foundation for the discussion of deployment implications, limitations, and future research in the following section.

Table 1: Comparative Evaluation of Conventional Monitoring and the Proposed Risk-Centric Methodology

Threat scenario	**Conventional rule-based monitoring**	**Proposed risk-centric methodology**
Unauthorized access	Detects predefined login violations	Correlates login anomalies with subsequent activity
Privilege misuse	Identifies configured authorization violations	Assesses privilege changes alongside user behavior and asset importance
Unusual transaction activity	Applies fixed transaction rules and thresholds	Identifies deviations from established transaction patterns
Suspicious event sequences	Detects predefined event combinations	Correlates related events across SAP and infrastructure logs
Alert prioritization	Uses predefined rule severity	Considers event severity, anomaly magnitude, and business impact

7. Discussion, Practical Implications, and Research Limitations

The proposed risk-centric methodology addresses a persistent challenge in SAP security monitoring: the difficulty of interpreting fragmented security events in a way that reflects both technical evidence and

business risk. Its principal contribution lies in combining security event correlation, statistical anomaly profiling, and risk-based prioritization within a common analytical framework. Conventional monitoring remains valuable for detecting recognizable violations, but isolated alerts may provide insufficient context when suspicious activity spans multiple applications, accounts, or infrastructure components. The proposed approach seeks to establish relationships among these observations and evaluate their significance against established behavioral patterns. This design offers a structured basis for investigating potential threats rather than relying exclusively on individual alerts. Its contribution should nevertheless be understood as a methodological framework whose practical effectiveness depends on validation, implementation quality, and the characteristics of the SAP environment under examination.

The integration of SAP-specific security evidence with enterprise SIEM analytics provides an important operational advantage. Application-level records can reveal transaction activity, authorization changes, and security-relevant operations, while infrastructure and identity records provide additional information about access paths, system interactions, and authentication behavior. Examining these sources together can help investigators reconstruct the sequence of events surrounding a suspicious incident. For example, an unusual login may become more significant when followed by an unexpected privilege change and access to a sensitive business function. However, correlation is useful only when the associated records are sufficiently reliable and the relationships among them are defensible. The methodology therefore emphasizes normalized event fields, consistent timestamps, contextual enrichment, and traceability to original records. Integration should improve the quality of investigation, not merely increase the volume of information collected.

Security analysts are among the primary beneficiaries of a risk-centric approach because their work requires distinguishing genuine threats from routine operational exceptions. Correlated event sequences can provide a clearer investigative starting point than numerous unrelated notifications, while contextual information can explain why a particular event has received elevated priority. Risk managers can use the resulting assessments to understand potential exposure involving critical business systems, sensitive information, and privileged accounts. Auditors may benefit from traceable event histories, documented detection criteria, and evidence supporting the investigation of security-relevant activities. Nevertheless, the framework should not replace professional judgment or established incident-handling procedures. Risk scores are analytical aids, not definitive determinations of malicious intent. Their usefulness depends on transparent scoring criteria, appropriate escalation rules, and the availability of supporting evidence that allows decisions to be independently reviewed.

Deployment across established SAP environments presents several practical challenges. Older application systems may have limited logging capabilities, inconsistent configurations, or dependencies that complicate the collection of security-relevant information. Different SAP components and supporting platforms may use dissimilar event structures, identifiers, severity classifications, and timestamp conventions. Normalizing these records requires care because removing source-specific attributes can eliminate information needed to establish the meaning of an event. In addition, increased collection and retention may create storage, processing, and administrative demands. Organizations should therefore begin by identifying critical assets, relevant audit sources, existing monitoring coverage, and the business processes most exposed to security risk. A phased implementation can establish reliable collection and normalization before introducing more extensive statistical profiling and cross-system correlation. This approach reduces deployment complexity while allowing detection quality to be assessed at each stage.

Baseline drift and statistical uncertainty introduce further limitations. Behavioral profiles represent observed activity over a particular reference period, but legitimate operating patterns change as responsibilities, transaction volumes, system configurations, and business schedules evolve. A baseline that accurately reflects one period may become inappropriate after a substantial organizational or technical change. Conversely, gradual malicious activity may remain close to established behavior and escape

detection if the analytical method is insufficiently sensitive. Sparse event data, incomplete records, and inconsistent user identification can also weaken statistical estimates. The methodology cannot eliminate these limitations merely by combining several analytical techniques. It requires periodic baseline review, documented threshold selection, explicit treatment of missing data, and continued examination of false-positive and false-negative cases. Where the available evidence is insufficient, the system should preserve uncertainty rather than present an unsupported conclusion.

Another limitation concerns the relationship between detection performance and operational value. A method may improve recall while producing enough additional alerts to reduce investigation efficiency. Similarly, stronger correlation may help reconstruct incidents but increase processing demands or introduce incorrect associations when timestamps and identifiers are unreliable. Risk scoring also involves judgment because event severity, asset importance, and potential business impact may be difficult to quantify consistently across different systems. Evaluation results must therefore be interpreted within the conditions under which they were obtained. Controlled test scenarios can support repeatability, but they may not reproduce every characteristic of production activity, including incomplete audit coverage, unusual business exceptions, and previously unseen attack sequences. The proposed methodology should consequently be regarded as a structured means of improving detection and investigation, not as a guarantee of comprehensive threat identification.

Future research should validate the methodology across different SAP configurations, business sectors, and operational scales. Comparative studies could examine whether the same correlation rules and statistical techniques remain effective across environments with different transaction volumes, user roles, and logging practices. Further experiments could assess the contribution of individual components through ablation testing, comparing results with and without contextual enrichment, statistical profiling, or cross-source correlation. Researchers should also investigate methods for adapting behavioral baselines while limiting the risk that malicious activity becomes incorporated into normal reference patterns. Additional work could examine the reproducibility of risk scores, the reliability of event association under imperfect timestamps, and the effects of incomplete data on detection performance. Such investigations would help establish which aspects of the framework generalize reliably and which require environment-specific configuration.

The broader implication is that SAP threat detection should be evaluated through the combined quality of its evidence, analytical decisions, and operational outcomes. Security event correlation establishes relationships among observations, statistical anomaly profiling identifies departures from established behavior, and risk assessment helps determine which findings deserve investigative attention. Their combination offers a coherent basis for examining threats that may not be apparent through isolated monitoring. However, the practical value of this approach depends on reliable source data, defensible analytical assumptions, and disciplined investigation procedures. By explicitly recognizing deployment constraints and statistical limitations, the methodology provides a foundation for further empirical study without overstating its demonstrated capabilities. Its long-term contribution will depend on whether subsequent evaluations establish repeatable improvements in detection quality, investigative efficiency, and risk-informed security operations across diverse SAP landscapes.

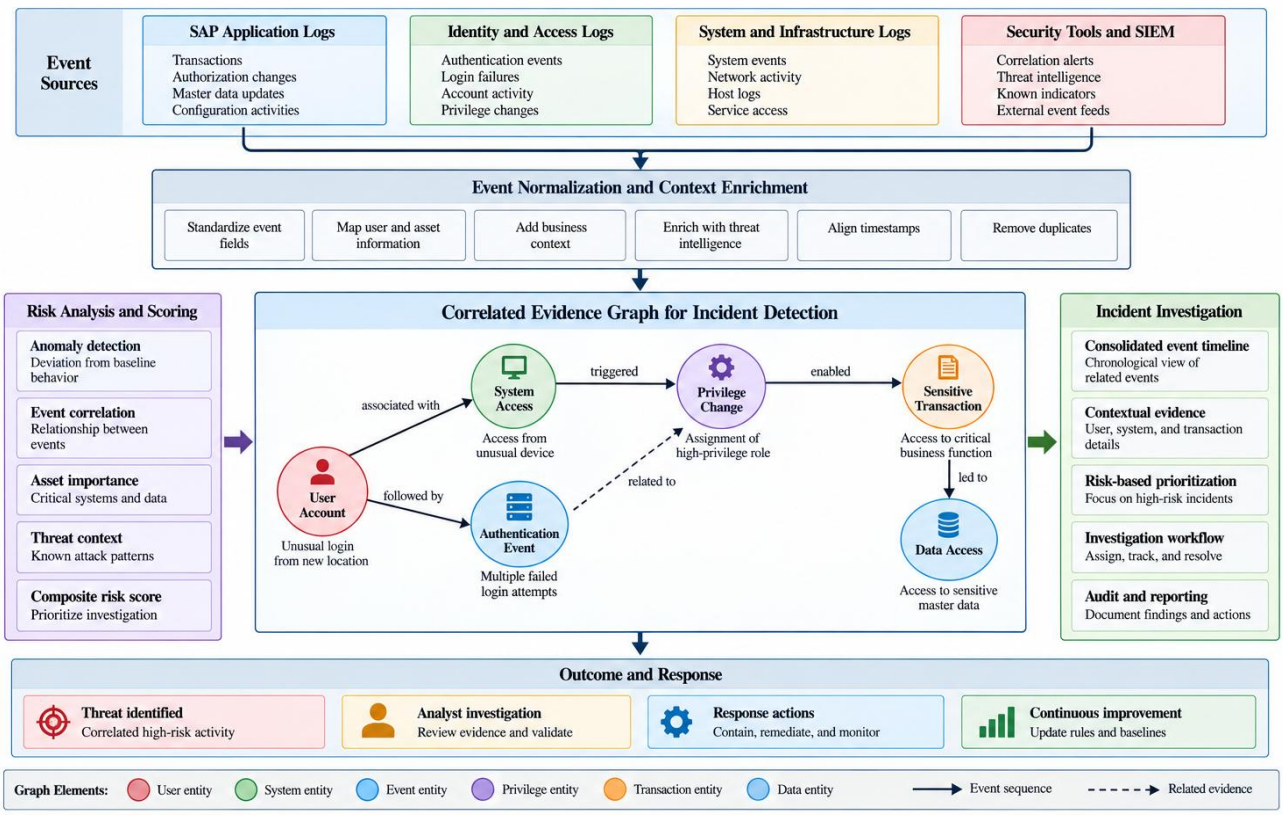


Figure 7: Operational Deployment and Incident Investigation Model

8. Conclusion and Future Work

The increasing complexity of interconnected SAP landscapes presents substantial challenges for identifying cyber threats through isolated log monitoring and conventional rule-based detection. Security-relevant activities may be distributed across application records, authentication histories, audit trails, transaction logs, and supporting infrastructure, making it difficult to reconstruct suspicious activity from individual alerts. This study addressed these challenges by proposing a risk-centric methodology that combines security event correlation, statistical anomaly profiling, and enterprise-wide Security Information and Event Management (SIEM) analytics. The primary objective was to establish a structured approach for identifying potentially malicious behavior, interpreting its operational context, and prioritizing investigations according to security significance and potential business impact.

The proposed methodology integrates six principal capabilities: event collection, data normalization, contextual enrichment, security event correlation, statistical anomaly assessment, and risk-based investigation. These capabilities establish a connected analytical process in which SAP-specific evidence can be examined alongside relevant infrastructure, identity, and system events. Behavioral baselines provide a reference for identifying deviations in user activity, transaction frequency, and administrative operations, while correlation mechanisms associate related observations across available sources. Risk assessment considers event severity, anomaly magnitude, asset importance, and potential business consequences to support investigative prioritization. This combination provides a methodological basis for examining suspicious activity more comprehensively than approaches that depend exclusively on isolated events or predefined detection conditions.

The evaluation framework establishes how the proposed methodology should be compared with conventional rule-based monitoring through unauthorized access, privilege misuse, unusual transaction activity, and suspicious event sequence scenarios. Precision, recall, false-positive rate, detection latency, correlation effectiveness, and investigation efficiency provide complementary measures for examining detection quality and operational usefulness. However, because measured experimental results have not been established in the material presented, the methodology's superiority cannot yet be claimed as an empirical finding. Its expected benefits, including improved contextual interpretation and more consistent alert prioritization, remain hypotheses requiring validation against documented test data. This distinction is essential to maintaining the credibility and reproducibility of the research.

The academic contribution lies in bringing SAP-specific security context, statistical anomaly profiling, and risk-based event correlation into a common analytical framework. Rather than treating anomaly detection as an independent classification task, the proposed approach examines deviations alongside related events, user responsibilities, affected assets, and business processes. This perspective provides a foundation for future studies investigating how application-level evidence can improve enterprise threat monitoring. The framework also establishes evaluation requirements that enable researchers to compare detection approaches using consistent definitions and transparent assumptions. Its broader research value will depend on whether subsequent experiments demonstrate repeatable improvements across different threat scenarios and SAP configurations.

From an operational perspective, the methodology offers a structured basis for connecting SAP Enterprise Threat Detection capabilities with enterprise SIEM workflows. Security analysts may use correlated evidence to reconstruct suspicious activity, while auditors can examine traceable event histories and documented investigation decisions. Enterprise risk managers may benefit from prioritization methods that consider the potential consequences of security incidents affecting critical business functions. Nevertheless, effective deployment requires reliable audit configuration, consistent timestamps, appropriate identity mapping, and sufficient logging coverage. Organizations must also account for integration costs, storage requirements, processing overhead, and the maintenance of detection rules and behavioral baselines. These practical considerations determine whether the proposed analytical capabilities can be sustained within established security operations.

Several limitations remain relevant to future investigation. Baseline drift, incomplete records, inconsistent event formats, and uncertain statistical estimates can affect detection reliability. Controlled or synthetic test scenarios may also fail to represent the full complexity of production environments. Further research should therefore evaluate the methodology using representative operational records, independently reviewed event labels, and repeatable comparative experiments. Additional studies could examine the contribution of individual detection components, the reliability of risk scores, and the effects of changing business activity on false-positive and missed-detection rates. Validation across different SAP landscapes would help determine which elements can be generalized and which require environment-specific adaptation.

The proposed risk-centric methodology establishes a coherent foundation for investigating cyber threats through the combined analysis of event relationships, behavioral deviations, and business risk. Its principal value is the structured connection between SAP security evidence and enterprise-wide analytical processes, supported by explicit evaluation criteria and recognition of operational limitations. Although its effectiveness must be established through empirical testing, the framework provides a practical and academically relevant basis for further research into SAP threat detection. Continued validation can clarify its contribution to security monitoring, incident investigation, and risk-informed decision-making, ultimately supporting more reliable protection of critical enterprise applications and business information.

9. References

[1] D'Antonio, S., Coppolino, L., et al. (2016). A framework for mastering heterogeneity in multi-layer security information and event correlation. Journal of Systems Architecture, 62, 78–88. https://doi.org/10.1016/j.sysarc.2015.11.010

[2] Granadillo, G. G., El-Barbori, M., et al. (2016). New types of alert correlation for security information and event management systems. In 2016 8th IFIP International Conference on New Technologies, Mobility and Security (NTMS). IEEE. https://doi.org/10.1109/NTMS.2016.7792462

[3] Kenaza, T., & Aiash, M. (2016). Toward an efficient ontology-based event correlation in SIEM. Procedia Computer Science, 83, 139–146. https://doi.org/10.1016/j.procs.2016.04.109

[4] Aburomman, A. A., & Reaz, M. B. I. (2017). A survey of intrusion detection systems based on ensemble and hybrid classifiers. Computers & Security, 65, 135–152. https://doi.org/10.1016/j.cose.2016.11.004

[5] Callegari, C., Giordano, S., et al. (2017). An information-theoretic method for the detection of anomalies in network traffic. Computers & Security, 70, 351–365. https://doi.org/10.1016/j.cose.2017.07.004

[6] Kamarudin, M. H., Maple, C., et al. (2017). A new unified intrusion anomaly detection in identifying unseen web attacks. Security and Communication Networks, 2017, Article 2539034. https://doi.org/10.1155/2017/2539034

[7] Viegas, E. K., Santin, A. O., et al. (2017). Toward a reliable anomaly-based intrusion detection in real-world environments. Computer Networks, 127, 200–216. https://doi.org/10.1016/j.comnet.2017.08.013

[8] Landauer, M., Wurzenberger, M., et al. (2018). Dynamic log file analysis: An unsupervised cluster evolution approach for anomaly detection. Computers & Security, 79, 94–116. https://doi.org/10.1016/j.cose.2018.08.009

[9] Khraisat, A., Gondal, I., et al. (2019). Survey of intrusion detection systems: Techniques, datasets and challenges. Cybersecurity, 2, Article 20. https://doi.org/10.1186/s42400-019-0038-7

[10] Gutierrez, R. J., Bauer, K. W., et al. (2018). Cyber anomaly detection: Using tabulated vectors and embedded analytics for efficient data mining. Journal of Algorithms & Computational Technology, 12(4), 293–310. https://doi.org/10.1177/1748301818791503

[11] Sapegin, A., Jaeger, D., et al. (2017). Towards a system for complex analysis of security events in large-scale networks. Computers & Security, 67, 16–34. https://doi.org/10.1016/j.cose.2017.02.001

[12] Du, M., Li, F., et al. (2017). DeepLog: Anomaly detection and diagnosis from system logs through deep learning. In Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, pp. 1285–1298. Association for Computing Machinery. https://doi.org/10.1145/3133956.3134015

[13] Boero, L., Cello, M., et al. (2017). Statistical fingerprint-based intrusion detection system (SF-IDS). International Journal of Communication Systems, 30(10), Article e3225. https://doi.org/10.1002/dac.3225

[14] Kabir, E., Hu, J., et al. (2018). A novel statistical technique for intrusion detection systems. Future Generation Computer Systems, 79(1), 303–318. https://doi.org/10.1016/j.future.2017.01.029

[15] Meng, W., Liu, Y., et al. (2019). LogAnomaly: Unsupervised detection of sequential and quantitative anomalies in unstructured logs. In Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, pp. 4739–4745. International Joint Conferences on Artificial Intelligence Organization. https://doi.org/10.24963/ijcai.2019/658

[16] Brown, A., Hutchinson, B., et al. (2018). Recurrent neural network attention mechanisms for interpretable system log anomaly detection. In Proceedings of the 1st Workshop on Machine Learning for Computing Systems. Association for Computing Machinery. https://doi.org/10.1145/3217871.3217872

[17] Liu, S., Chen, X., et al. (2019). Network log anomaly detection based on GRU and SVDD. In 2019 IEEE International Conference on Parallel & Distributed Processing with Applications, Big Data & Cloud Computing, Sustainable Computing & Communications, Social Computing & Networking, pp. 1244–1249. IEEE. https://doi.org/10.1109/ISPA-BDCloud-SustainCom-SocialCom48970.2019.00177

[18] He, P., Zhu, J., et al. (2016). An evaluation study on log parsing and its use in log mining. In 2016 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, pp. 654–661. IEEE. https://doi.org/10.1109/DSN.2016.66

[19] Hamooni, H., Debnath, B., et al. (2016). LogMine: Fast pattern recognition for log analytics. In Proceedings of the 25th ACM International Conference on Information and Knowledge Management, pp. 1573–1582. Association for Computing Machinery. https://doi.org/10.1145/2983323.2983358

[20] Zhu, H., & Liu, W. (2017). A universal high-performance correlation analysis detection model and algorithm for network intrusion detection system. Mathematical Problems in Engineering, 2017, Article 8439706. https://doi.org/10.1155/2017/8439706

[21] Barzegar, M., & Shajari, M. (2018). Attack scenario reconstruction using intrusion semantics. Expert Systems with Applications, 108, 119–133. https://doi.org/10.1016/j.eswa.2018.04.030

[22] Daneshgar, F. F. (2016). Extracting fuzzy attack patterns using an online fuzzy adaptive alert correlation framework. Security and Communication Networks, 9(14), 2245–2260. https://doi.org/10.1002/sec.1483

[23] Salah, S., Maciá-Fernández, G., et al. (2019). Fusing information from tickets and alerts to improve the incident resolution process. Information Fusion, 45, 38–52. https://doi.org/10.1016/j.inffus.2018.01.011

[24] Gallos, L. K., Korczyński, M., et al. (2017). Anomaly detection through information sharing under different topologies. EURASIP Journal on Information Security, 2017, Article 5. https://doi.org/10.1186/s13635-017-0056-5

[25] Cordy, M., Muller, S., et al. (2019). Search-based test and improvement of machine-learning-based anomaly detection systems. In Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis. Association for Computing Machinery. https://doi.org/10.1145/3293882.3330580

[26] Rajivan, P., & Cooke, N. J. (2018). Information-pooling bias in collaborative security incident correlation analysis. Human Factors, 60(5), 626–639. https://doi.org/10.1177/0018720818769249

[27] Qin, T., Gao, Y., et al. (2018). Potential threats mining methods based on correlation analysis of multi-type logs. IET Networks, 7(5), 299–305. https://doi.org/10.1049/iet-net.2017.0188

[28] Aghaei, E., & Serpen, G. (2017). Ensemble classifier for misuse detection using N-gram feature vectors through operating system call traces. Journal of Intelligent & Fuzzy Systems, 33(3), 1747–1761. https://doi.org/10.3233/HIS-170247

[29] Luh, R., Marschalek, S., et al. (2017). Semantics-aware detection of targeted attacks: A survey. Journal of Computer Virology and Hacking Techniques, 13, 47–85. https://doi.org/10.1007/s11416-016-0273-3

[30] Shu, X., Yao, D., et al. (2017). Long-span program behavior modeling and attack detection. ACM Transactions on Privacy and Security, 20(4), Article 12. https://doi.org/10.1145/3105761