

SIEM Workflow Optimization and Threat Intelligence for Cloud Incident Response: Approaches to Correlation Automation, Triage Reduction, and Cloud Infrastructure Resilience

 **Kateryna Oblakevych**

Centific, Redmond, WA - Software Engineer
Seattle, WA, USA

RECEIVED - 02-03-2026, RECEIVED REVISED VERSION - 16-03-2026, ACCEPTED- 22-03-2026, PUBLISHED- 07-04-2026

Abstract

Cloud infrastructure faces a growing volume of security events that legacy rule-based Security Information and Event Management (SIEM) systems struggle to process in time for effective incident response. According to the IBM Cost of a Data Breach Report 2024, the global average breach cost reached USD 4.88 million, with multi-cloud environments accounting for 40 percent of affected organizations and breach lifecycles averaging 258 days. This article examines approaches to SIEM workflow optimization in cloud environments, with a focus on automated correlation, behavior-based anomaly scoring, and threat intelligence integration as means to reduce incident triage time and improve response quality. The methodology combines a systematic literature review of peer-reviewed publications from 2019 to 2025, content analysis of industry benchmark reports, and synthesis of practical engineering experience from cloud data security initiatives. Results show that ML-augmented SIEM architectures can reduce triage time by 45 to 78 percent compared to manual processes, lower false-positive rates significantly, and support automated containment of priority-one incidents. The study proposes an integrated model linking Zero Trust access controls, Kubernetes audit logging, HashiCorp Vault secrets management, and MITRE ATT&CK-enriched correlation into a coherent cloud incident response architecture. The findings are relevant to cloud security engineers, security operations center analysts, and organizations building or modernizing security operations for critical cloud systems.

Keywords: SIEM, cloud incident response, anomaly detection, threat intelligence, Zero Trust, DevSecOps, MITRE ATT&CK, Kubernetes security, SOAR automation, data exfiltration prevention.

Introduction The adoption of cloud-native architectures has accelerated significantly across the enterprise sector during the period 2022 to 2025, bringing with it an expanded attack surface that traditional security information and event management (SIEM) systems were not designed to address at scale. Organizations operating in multi-cloud and hybrid environments generate log volumes that can exceed tens of billions of events per day, creating conditions in which security operations center (SOC) analysts face alert fatigue, delayed triage, and inconsistent incident

classification [3]. According to the IBM Cost of a Data Breach Report 2024, 40 percent of breaches involved data stored across multiple cloud environments, and those breaches cost organizations an average of USD 5.17 million, while taking the longest time to identify [1]. These figures indicate that the gap between breach occurrence and detection remains one of the most costly problems in enterprise security.

Research published in ACM Computing Surveys confirms that SOC analysts receive alert volumes so large that false-positive rates in some environments approach 99 percent with rule-based correlation, creating a situation in which genuine threats are buried within noise [2]. The same study identifies alert triage as the primary bottleneck in incident response workflows. Parallel findings from IBM indicate that organizations using AI and automation extensively in their security operations identified and contained breaches approximately 100 days faster than those relying on manual processes, and incurred average breach costs of USD 3.84 million, compared to USD 5.72 million for organizations without such technologies [1]. These data points provide quantitative grounding for the central argument of this article: that SIEM workflow automation, when combined with behavior-based anomaly scoring and threat intelligence enrichment, produces measurable improvements in cloud incident response outcomes.

The security challenges specific to cloud infrastructure differ from those in traditional on-premises environments in several respects. Kubernetes-orchestrated workloads are ephemeral, making persistent identity and access controls more complex. CI/CD pipelines introduce automated deployment paths that, without proper security gates, can propagate vulnerable or misconfigured containers directly into production. Secrets management at scale, particularly across multiple deployment environments, presents an additional risk when static credentials or improperly protected environment variables are used instead of dynamic secret issuance through systems such as HashiCorp Vault [7]. These architectural characteristics require a SIEM strategy that goes beyond passive log aggregation toward active, context-aware correlation tied to the specific control structures of cloud-native environments.

The scientific gap addressed by this article is the absence of an integrated engineering framework that combines SIEM correlation automation, behavior-based anomaly scoring, and cloud-specific controls (Zero Trust, RBAC, Kubernetes audit, secrets management) into a unified incident response architecture, validated through both empirical literature and practical implementation experience.

The goal of this study is to analyze existing approaches to SIEM workflow optimization in cloud environments, propose an integrated architecture for automated cloud incident response, and evaluate its

constituent elements against published performance benchmarks from peer-reviewed sources and industry reports.

Scientific novelty: This article introduces an author-developed composite model that integrates behavior-based anomaly scoring with MITRE ATT&CK-enriched correlation and cloud-native security controls into a structured cloud incident response workflow, offering a synthesis that has not been previously presented in this specific configuration.

The hypothesis of this research is that combining ML-augmented SIEM correlation with Zero Trust access controls, Kubernetes audit logging, and dynamic secrets management produces incident response outcomes that are measurably superior to rule-based SIEM systems operating in isolation from the broader cloud security architecture.

Limitations of this study include its reliance on published benchmark data and practitioner-reported metrics rather than controlled laboratory experiments, which constrains the precision of comparative performance claims. The paper does not produce or evaluate custom software implementations; it synthesizes and analyzes existing approaches. Additionally, performance figures from specific vendor platforms may not generalize uniformly across all cloud environments or organizational maturity levels. These constraints are typical of systematic review and comparative analysis methodologies.

The practical value of this work extends to the United States technology sector specifically. As cloud-native security engineering becomes a critical competency for organizations managing regulated data and critical infrastructure, the ability to design and deploy automated SIEM architectures that reduce breach detection time and contain data exfiltration at scale directly supports national cybersecurity resilience objectives. Engineers and security architects operating at the intersection of cloud platform development, DevSecOps automation, and threat intelligence can draw on this framework to accelerate the security maturity of systems that process sensitive data for enterprise and government clients.

Materials and Methods

The research methodology combines three complementary approaches: systematic literature review,

content analysis of industry benchmark reports, and synthesis of practical engineering experience from cloud security implementations.

A structured literature search was conducted across peer-reviewed databases including ACM Digital Library, IEEE Xplore, SpringerLink, and ScienceDirect. Search terms included combinations of "SIEM," "cloud incident response," "anomaly detection," "SOAR automation," "Zero Trust," "Kubernetes security," "threat intelligence," "DevSecOps," and "behavior-based scoring." The search was limited to publications from 2019 to 2025, yielding an initial set of 87 candidate publications. Following title and abstract screening, full-text review, and exclusion of sources that did not address cloud-specific security operations or SIEM functionality, 16 peer-reviewed articles and conference papers were included in the final corpus. Selection prioritized publications indexed in Scopus or the Web of Science that reported empirical results, comparative evaluations, or systematic analysis of SIEM-related security outcomes.

Two industry benchmark reports were incorporated as a secondary source tier: the IBM Cost of a Data Breach Report 2024 (produced by the Ponemon Institute on behalf of IBM and based on analysis of 604 breached organizations across 17 industries in 16 countries) and the CrowdStrike Fal.Con 2024 technical analysis of next-generation SIEM automation. These reports were selected for their methodological transparency and breadth of empirical coverage. Their data were used specifically to provide quantitative context for breach cost, detection timelines, and automation impact metrics, rather than as primary theoretical sources.

This study incorporates a structured synthesis of engineering experience from cloud data security initiatives in which the author participated directly. The relevant projects include (a) a cloud data security initiative focused on protecting enterprise cloud infrastructure from data exfiltration and insider threats, in which SIEM correlation workflows were automated, resulting in an approximately 45 percent reduction in incident triage time; (b) the development and deployment of a behavior-based anomaly scoring approach for security event prioritization; (c) the automation of Kubernetes-based deployments

incorporating security controls such as OPA policy enforcement and image scanning, reducing manual release effort by approximately 80 percent; and (d) the integration of HashiCorp Vault and Consul for dynamic secrets management across multiple deployment environments. These practical reference points are used to ground the proposed framework in operational reality rather than solely in theoretical recommendation.

Results and Discussion

Security information and event management systems occupy the operational center of most enterprise security architectures, aggregating logs from endpoints, network devices, cloud services, identity providers, and application layers, then applying correlation rules to produce alerts [14]. In cloud environments, the source diversity and event velocity are substantially higher than in traditional on-premises deployments: a single Kubernetes cluster running 50 to 100 microservices can produce several million audit log entries per day, and multi-cloud environments compound this volume across AWS, GCP, and Azure control planes simultaneously.

Rule-based correlation engines, which dominated SIEM architectures throughout the 2010s, apply fixed logical conditions to incoming event streams. While effective for detecting well-defined attack patterns such as specific malware signatures or known indicators of compromise (IoCs), rule-based systems struggle with novel attack behavior, low-and-slow data exfiltration, and multi-stage campaigns that span legitimate-looking API calls across extended timeframes [16]. Alahmadi, Axon, and Martinovic [2] documented through qualitative interviews with SOC analysts that in environments using purely rule-based SIEM, analysts reported false-positive rates approaching 99 percent, with the practical consequence that analysts adopt informal triage shortcuts that can cause genuine threats to be dismissed along with noise.

The table below compares performance indicators across three SIEM deployment configurations, based on the IBM 2024 report and published ML-SIEM research. These figures provide a quantitative basis for evaluating the relative benefit of successive levels of automation.

Table 1. Comparative performance indicators for SIEM configurations by automation level (compiled by the author based on [1, 2, 4, 7, 11, 14]).

Metric	Legacy SIEM(rule-based only)	SIEM + ML anomaly scoring (partial)	SIEM + ML +automated triage
Mean time to identify breach (days)	194	220-240	177
Mean time to contain breach (days)	64	65-75	51
Average breach cost (USD million)	5.72	4.64	3.84
SOC analyst alert false-positive rate (%)	~99%	~60%	~40-50%
Automated response coverage (%)	<5%	30-40%	>60%
Triage time reduction vs. manual baseline	0%	~35%	~45-78%

The data in Table 1 show a clear gradient: adding ML-based anomaly scoring to rule-based correlation reduces the breach lifecycle and alert false-positive rate, while adding automated triage through SOAR integration produces the strongest combined effect on both time and cost metrics.

The 45 to 78 percent range for triage time reduction spans results from the IBM report (approximately 45 percent reduction in detection and containment time for extensive AI users versus non-users [1]) and the UEBA simulation study by Rabbani et al. [11].

Figure 1 below presents the average data breach cost stratified by AI and automation adoption level, drawn directly from the IBM 2024 report. This chart illustrates the financial incentive for organizations to move from manual processes toward automated security operations.




Figure 1. Average data breach cost by AI/automation adoption level (compiled by the author based on [1]).

Building on the literature review and practical experience synthesized in Section 2, this section presents an author-developed integrated architecture for automated cloud incident response. The architecture addresses four principal challenges identified in the literature: high alert volume and false-positive rates, slow detection-to-containment cycles, insufficient threat context for effective analyst decision-making, and lack of integration between SIEM correlation and the specific control structures of cloud-native infrastructure.

The proposed workflow is structured in five functional layers: (1) log ingestion and normalization, (2) SIEM correlation with ML anomaly scoring, (3) alert enrichment and prioritization, (4) automated response through SOAR playbooks, and (5) audit logging and compliance trail generation. Figure 2 presents this architecture as a process diagram.




Figure 2. Proposed SIEM workflow for automated cloud incident response (compiled by the author based on [1, 4, 5, 7, 11, 14]).

In the ingestion layer, log sources are normalized to a common schema (Common Event Format or Elastic Common Schema) before entering the correlation engine. Normalization is a prerequisite for cross-source correlation and is a step that many organizations treat as an afterthought, leading to correlation gaps. Practical experience from cloud data security initiatives confirms that inconsistent field naming between Kubernetes audit logs and cloud provider control plane events is a frequent source of missed detections.

The correlation engine applies both deterministic rules (for known-bad IoCs and compliance-critical events such as privilege escalation in production namespaces) and probabilistic ML-based scoring. The ML component draws on behavioral baselines established over rolling time windows to identify deviations in user session behavior, API call patterns, resource access rates, and data transfer volumes. Alert enrichment adds contextual information from MITRE ATT&CK mappings, threat intelligence feeds, and identity context derived from the RBAC graph before alerts reach the analyst queue. This enrichment step is critical: research by González-Granadillo, González-Zarzosa, and Diaz [14] found that alert context quality is the primary determinant of analyst decision accuracy and speed.

A core component of the proposed architecture is a behavior-based anomaly scoring model that assigns a composite risk score (0 to 100) to each security event based on deviations from established behavioral baselines. The model is designed as a modular overlay on top of existing SIEM infrastructure, consuming normalized events and producing scored records that can be ingested by standard SOAR platforms.

The scoring model uses five input signal categories: user session behavior (login time, location, session duration, action rate), API call patterns (endpoint frequency, request volume, authentication method used), resource access rate (files accessed per minute, data volume read), authentication events (OAuth2 token issuance, JWT claim deviations, RBAC role changes), and data transfer volume (egress rates relative to rolling 7-day baseline). Each signal category contributes a partial score using Z-score normalization against a sliding-window baseline, and the partial scores are combined using a weighted average calibrated during an initial learning period.

Figure 3 presents the anomaly scoring model as a data flow diagram, showing the progression from input signals through feature extraction, scoring, and threshold-based prioritization.




Figure 3. Behavior-based anomaly scoring model for cloud threat prioritization (compiled by the author based on [5, 9, 11, 16]).

The threshold boundaries in Figure 3 (auto-block at score above 80, analyst alert at 50 to 80, log and monitor at 20 to 50, discard below 20) reflect the prioritization logic applied in the practical cloud security initiative referenced in Section 2. These thresholds are not universal; they require calibration against organization-specific false-positive tolerance and staffing capacity. However, the threshold structure itself, tiered by automated action type, reflects a design principle supported in the ACM Computing Surveys review of alert fatigue mitigation: automated disposition of the lowest-risk tier and the highest-confidence high-risk tier significantly reduces the cognitive load on analysts without increasing the risk of undetected incidents [3].

The ML-augmented SIEM study by Rabbani et al. [11] found that Transformer-GNN ensemble architectures applied to UEBA data achieved an F1-score of 0.90 and a 40 percent

reduction in false positives compared to rule-based correlation. Those results align directionally with the behavior-based scoring approach described here, and provide independent quantitative support for the architecture proposed in this article.

One of the distinctive contributions of this article is the explicit mapping of SIEM integration points to specific cloud-native security controls. Most published SIEM research treats cloud infrastructure as a generic log source. In practice, cloud-native environments have distinct control structures (RBAC policies, Kubernetes audit APIs, service mesh mTLS configurations, Vault audit devices) that each produce security-relevant events and each represent specific attack surfaces. The table below maps these control layers to their threat profiles and SIEM integration points.

Table 2. Cloud-native security control layers and SIEM integration points (compiled by the author based on [7, 8, 12, 14, 15, 18]).

Security layer	Technology / standard	SIEM integration point	Threat addressed
Identity & access	OAuth2, JWT, RBAC, MFA	Auth event logs	Credential abuse, unauthorized lateral movement
Container orchestration	Kubernetes audit, OPA policies	Audit log stream to SIEM	Privileged pod creation, policy bypass

Secrets management	HashiCorp Vault, Consul	Vault audit device	Secret leakage, insecure credential storage
CI/CD pipeline	SAST, DAST, SCA, image scanning	Build artifact events	Supply chain compromise, vulnerable images
Data protection	PII encryption, GDPR logging controls	DLP event alerts	Data exfiltration, insider threat
Network & API	mTLS, Zero Trust segmentation, WAF	Network flow anomalies	API abuse, east-west lateral movement

The secrets management row in Table 2 reflects a specific risk pattern observed in multi-environment cloud deployments: when static credentials are stored in environment variables or Kubernetes native secrets (which are stored as base64-encoded data in etcd), they are accessible to any process with sufficient RBAC permissions and are not automatically rotated. HashiCorp Vault addresses this through dynamic short-lived credentials with automatic revocation, and Vault audit device produces a complete log of every secret access event that can be streamed to the SIEM for anomaly detection [7]. Integration of Vault audit logs with behavioral scoring enables detection of unusual access patterns such as a service account requesting secrets outside its normal operating hours or from an unexpected IP range.

Kubernetes audit logging provides a different type of signal: a record of every API server interaction, including pod

creation, role binding changes, secret reads, and namespace modifications. Mapping these events to MITRE ATT&CK cloud-specific techniques (particularly the Discovery, Lateral Movement, and Persistence tactic categories) enables detection of multi-stage attacks that would be invisible to network-layer monitoring alone. Research by Al-Sada et al. [12] confirms that ATT&CK-aligned detection coverage increases the proportion of detected multi-step campaigns compared to indicator-based approaches.

The CI/CD pipeline is both a security control point and a data source for the SIEM. Every build event, container image scan result, static analysis finding, and deployment action produces a record that, when ingested into the SIEM, extends visibility to the software supply chain. Figure 4 shows the proposed secure CI/CD pipeline architecture with integrated SIEM audit logging.




Figure 4. Secure CI/CD pipeline with integrated SIEM audit logging and security gates (compiled by the author based on [6, 7, 8, 10, 15, 18]).

The pipeline in Figure 4 includes static analysis and software composition analysis at the pull request stage (blocking deployment of code with high-severity vulnerabilities), container image hardening and scanning in the build stage, dynamic application security testing against a staging environment, and a secrets scan using Vault integration to ensure no credentials are embedded in build artifacts. All gate outcomes are logged to the SIEM, creating an audit trail that satisfies both operational security requirements and compliance frameworks such as SOC 2 Type II. The continuous SIEM monitoring bar at the base of the diagram represents the streaming of pipeline events to the correlation engine in real time, so that anomalous build behavior (such as a previously unseen image being pushed to a production registry, or a build job running outside scheduled hours) triggers an alert.

The practical Kubernetes deployment automation experience referenced in Section 2 produced approximately an 80 percent reduction in manual release work through standardized Helm-based deployments with OPA policy checks embedded as admission controller webhooks. This reduction also improved security consistency: when deployments are automated rather than manually executed, the probability of security control misconfiguration is lower because policy enforcement occurs programmatically at every deployment rather than depending on individual analyst adherence to checklist procedures [8, 13].

Table 3 synthesizes the performance outcomes associated with each element of the proposed architecture, drawing on both published research and practical implementation data. It serves as a reference point for organizations planning phased adoption of the integrated model.

Table 3. Performance benchmarks for integrated SIEM workflow elements (compiled by the author based on [1, 4, 7,14, 16,17, 19]).

Implementation element	Target metric / outcome	Baseline before optimization	Expected outcome after full deployment
ML-based anomaly scoring layer	False-positive rate	~99% (rule-based)	~40-50% reduction vs. baseline
Automated SIEM correlation workflows	Incident triage time	Manual; hours per incident	~45% reduction observed in practice; up to 78% reported
SOAR playbook automation	Automated response coverage	<5%	>60% of P1/P2 alerts auto-contained
Kubernetes CI/CD automation	Manual release effort	~100% manual	~80% reduction in manual work
Vault/Consul secrets management	Dynamic secret rotation	Static, manually rotated	Automatic rotation; full audit trail in SIEM
MITRE ATT&CK enrichment of alerts	Detection coverage	Ad-hoc rule matching	Structured TTP mapping; improved lateral movement detection

Several recommendations emerge from the synthesis in Table 3 and the broader analysis of the literature.

First, organizations should prioritize the ML anomaly scoring layer before expanding SOAR automation. The reduction in false-positive rate is a prerequisite for effective

automated response: automating responses to 99 percent false-positive alert streams creates operational disruption without security benefit. Calibration of the scoring model should use a minimum 30-day behavioral baseline before live scoring is enabled in production.

Second, Kubernetes audit log streaming should be configured from the initial cluster deployment, not added retrospectively. Retrospective enabling of audit logging in production clusters can cause latency spikes and requires careful tuning of the audit policy to avoid logging health check and high-frequency service account events that produce noise without security value. The recommended practice is to define audit policies aligned to the MITRE ATT&CK cloud techniques relevant to the organizations threat model [12] before the cluster is promoted to production.

Third, the secrets management and SIEM integration should be treated as a single project rather than two separate initiatives. Vault audit device output is only useful to the SIEM if the correlation rules are designed to consume it. This requires coordination between the platform engineering team and the security operations team, a boundary that in many organizations is organizationally siloed. Practical experience from multi-environment Vault and Consul deployments confirms that this coordination gap is the most common reason organizations have Vault deployed but not properly integrated into their detection stack.

Fourth, the MITRE ATT&CK mapping should be applied as a continuous enrichment process rather than a one-time rule configuration exercise. The ATT&CK framework is updated multiple times per year to reflect new adversary techniques observed in the wild. Leite et al. [6] demonstrated that automated extraction of ATT&CK techniques from threat intelligence reports and their integration into SIEM correlation rules allows the detection capability to remain current with evolving threat actor behavior, an advantage that static rule sets cannot replicate without significant manual maintenance effort.

Fifth, GDPR-aligned audit logging and PII encryption controls should be built into the SIEM data pipeline from the outset. Cloud incident response investigations frequently require access to log data that may contain personal information. Establishing privacy-by-design data governance before an incident occurs, rather than after, ensures that the investigation does not create a secondary compliance risk and that the audit trail itself is admissible in regulatory proceedings [7, 15, 20].

Conclusion

This article examined approaches to SIEM workflow optimization for cloud incident response, with a focus on ML-augmented correlation automation, behavior-based anomaly scoring, and integration with cloud-native security controls. The analysis demonstrates that the gap between breach occurrence and detection, which averaged 258 days across all organizations studied in the IBM 2024 benchmark, can be substantially reduced through systematic application of the architecture elements described in this article.

The principal findings are as follows. ML-based anomaly scoring, when applied on top of rule-based SIEM correlation, reduces false-positive rates from approximately 99 percent to 40 to 60 percent, according to both published research and practical implementation data. Organizations using AI and automation extensively in their security operations incur average breach costs approximately 33 percent lower than organizations without such technologies, per IBM 2024. Automated SIEM correlation workflows, as implemented in the author-referenced cloud data security initiative, reduced incident triage time by approximately 45 percent, consistent with the broader range of 45 to 78 percent reported in the ML-SIEM literature. The integration of Kubernetes audit logging, HashiCorp Vault audit events, and CI/CD pipeline security gate outputs into a unified SIEM correlation engine extends detection coverage to attack vectors that network-layer monitoring cannot observe. MITRE ATT&CK-enriched alert context improves analyst decision quality and reduces the time required to determine incident scope and priority.

The integrated architecture proposed in this article links these elements into a coherent workflow: normalized log ingestion, ML-augmented correlation, behavior-based anomaly scoring with tiered automated response, and continuous MITRE ATT&CK-aligned enrichment, supported by Zero Trust access controls and dynamic secrets management. This configuration addresses the scientific gap identified in the introduction: the absence of a published framework that integrates SIEM automation with the specific control structures of cloud-native infrastructure.

For security engineers and SOC architects working on cloud systems, the practical value of this framework lies in its phased implementability: the anomaly scoring model and Vault audit integration can be deployed

independently of each other, and each produces measurable value without requiring the full architecture to be in place simultaneously. This makes the framework applicable to organizations at different stages of security maturity, from those modernizing legacy SIEM deployments to those building cloud-native security operations from the ground up.

References

1. IBM. (2024). Cost of a data breach report 2024. Retrieved from: https://www-api.ibm.com/adobe/assets/urn%3Aaaid%3Aaem%3Ac4711149-da99-4654-9114-8ee524108540/original/as/cost_of_a_data_breach_report_2024.pdf (date accessed: March 12, 2026).
2. Alahmadi, B. A., Axon, L., & Martinovic, I. (2022). 99% false positives: A qualitative study of SOC analysts' perspectives on security alarms. In 31st USENIX Security Symposium (USENIX Security 22) (pp. 2783–2800). USENIX Association. Retrieved from: <https://www.usenix.org/conference/usenixsecurity22/presentation/alahmadi> (date accessed: April 9, 2026).
3. Tariq, S., Baruwat Chhetri, M., Nepal, S., & Paris, C. (2025). Alert fatigue in security operations centres: Research challenges and opportunities. *ACM Computing Surveys*, 57(9), Article 224, 1–38. <https://doi.org/10.1145/3723158>.
4. Bridges, R. A., Rice, A. E., Oesch, S., Nichols, J. A., Watson, C., Spakes, K., Norem, S., Huettel, M., Jewell, B., Weber, B., Gannon, C., Bizovi, O., Hollifield, S. C., & Erwin, S. (2023). Testing SOAR tools in use. *Computers & Security*, 129, 103201. <https://doi.org/10.1016/j.cose.2023.103201>.
5. Ban, T., Takahashi, T., Ndichu, S., & Inoue, D. (2023). Breaking alert fatigue: AI-assisted SIEM framework for effective incident response. *Applied Sciences*, 13(11), 6610. <https://doi.org/10.3390/app13116610>.
6. Leite, C., den Hartog, J., dos Santos, D. R., & Costante, E. (2022). Actionable cyber threat intelligence for automated incident response. In H. P. Reiser and M. Kyas (Eds.), *Secure IT Systems: 27th Nordic Conference, NordSec 2022* (Lecture Notes in Computer Science, Vol. 13700, pp. 368–385). Springer. https://doi.org/10.1007/978-3-031-22295-5_20.
7. Chandramouli, R., & Butcher, Z. (2023). A zero trust architecture model for access control in cloud-native applications in multi-cloud environments (NIST Special Publication 800-207A). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-207A>.
8. Naseer, A., Naseer, H., Ahmad, A., Maynard, S. B., & Siddiqui, A. M. (2023). Moving towards agile cybersecurity incident response: A case study exploring the enabling role of big data analytics-embedded dynamic capabilities. *Computers & Security*, 135, 103525. <https://doi.org/10.1016/j.cose.2023.103525>.
9. Guo, H., Yuan, S., & Wu, X. (2021). LogBERT: Log anomaly detection via BERT. In 2021 International Joint Conference on Neural Networks (IJCNN) (pp. 1–8). IEEE. <https://doi.org/10.1109/IJCNN52387.2021.9534113>.
10. Dhanapala, I. S. A., Bharti, S., McGibney, A., & Rea, S. (2024). Toward a performance-based trustworthy edge-cloud continuum. *IEEE Access*, 12, 99201–99212. <https://doi.org/10.1109/ACCESS.2024.3429197>.
11. Rabbani, M., Rashidi, L., & Ghorbani, A. A. (2024). A graph learning-based approach for lateral movement detection. *IEEE Transactions on Network and Service Management*, 21(5), 5361–5373. <https://doi.org/10.1109/TNSM.2024.3414267>.
12. Al-Sada, B., Sadighian, A., & Oligeri, G. (2024). MITRE ATT&CK: State of the art and way forward. *ACM Computing Surveys*, 57(1), Article 12, 1–37. <https://doi.org/10.1145/3687300>.
13. Baruwat Chhetri, M., Tariq, S., Singh, R., Jalalvand, F., Paris, C., & Nepal, S. (2024). Towards human-AI teaming to mitigate alert fatigue in security operations centres. *ACM Transactions on Internet Technology*, 24(3), Article 12, 1–22. <https://doi.org/10.1145/3670009>.
14. González-Granadillo, G., González-Zarzosa, S., & Diaz, R. (2021). Security information and event management (SIEM): Analysis, trends, and usage in critical infrastructures. *Sensors*, 21(14), 4759. <https://doi.org/10.3390/s21144759>.

15. Sarkar, S., Choudhary, G., Shandilya, S. K., Hussain, A., & Kim, H. (2022). Security of zero trust networks in cloud computing: A comparative review. *Sustainability*, 14(18), 11213. <https://doi.org/10.3390/su141811213>.
16. Tendikov, N., Rzaeva, L., Saoud, B., Shaya, I., Azmi, M. H., Myrzatay, A., & Alnakhli, M. (2024). Security information event management data acquisition and analysis methods with machine learning principles. *Results in Engineering*, 22, 102254. <https://doi.org/10.1016/j.rineng.2024.102254>.
17. Charla, R. R. (2025). AI-enhanced automated incident response in SIEM with explainability for SOC analysts. In 2025 20th International Joint Symposium on Artificial Intelligence and Natural Language Processing (iSAI-NLP) (pp. 1–10). IEEE. <https://doi.org/10.1109/iSAI-NLP66160.2025.11320514>.
18. Nelson, A., Rekhi, S., Souppaya, M., & Scarfone, K. A. (2025). Incident response recommendations and considerations for cybersecurity risk management: A CSF 2.0 community profile (NIST Special Publication 800-61 Rev. 3). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-61r3>.
19. Khan, M. A., Azim, A., Abazari, F., Eargle, F., & Gardiner, J. (2024). Automated offense prioritization for SIEM using probabilistic machine learning models. In *Proceedings of the 37th Canadian Conference on Artificial Intelligence (Canadian AI 2024)*. <https://doi.org/10.21428/594757db.733c1367>.
20. Amro, A., Gkioulos, V., & Katsikas, S. K. (2023). Assessing cyber risk in cyber-physical systems using the ATT&CK framework. *ACM Transactions on Privacy and Security*, 26(2), Article 22, 1–33. <https://doi.org/10.1145/3571733>.