Review of Computer Engineering Research
2026 Vol. 13, No. 3, pp. 227-243
ISSN(e): 2410-9142
ISSN(p): 2412-4281
DOI: 10.18488/

A rule-driven SOC architecture for real-time threat detection and autonomous incident response

Jilika Jithendarnadh[1+]
J Balaraju[2]

[1,2]*Department of Computer Science Engineering, School of Engineering Anurag University, Hyderabad, India.*
[1]*Email: jithendarnadh@gmail.com*
[2]*Email: jb7443@gmail.com*

(+ Corresponding author)

Article History
Received: 2 March 2026
Revised: 29 May 2026
Accepted: 18 June 2026
Published: 4 September 2026

Keywords
Cloud security telemetry
Incident response automation
MITRE ATT&CK
Rule-based security analytics
Security operations center
Threat detection.

ABSTRACT

Security Operations Centers (SOCs) are essential for monitoring and responding to cyber threats in cloud-native environments, where infrastructure is dynamic, multi-tenant, and API-driven. Conventional SOCs rely heavily on manual triage and SIEM-based alerting, resulting in delayed detection of cloud-specific attacks such as IAM misuse, privilege escalation, and data exfiltration. This paper proposes a rule-driven SOC architecture that integrates cloud telemetry, threat intelligence, MITRE ATT&CK Cloud Matrix mappings, and SOAR-based response automation. The framework normalizes heterogeneous logs, correlates multi-stage attack behaviors, enriches alerts with threat intelligence, and executes automated or semi-automated remediation actions. Evaluation using AWS CloudTrail, VPC Flow Logs, Azure Activity Logs, and Kubernetes audit logs demonstrates a 41.2% reduction in time-to-detection, a 37.5% reduction in time-to-response, and 87.9% precision in automated remediation.

Contribution/Originality: This study contributes to cloud SOC literature by proposing an ATT&CK-aligned, rule-driven architecture that combines stateful multi-stage correlation, policy-aware autonomous response, and continuous feedback refinement. This study uses a formal rule model linking cloud telemetry, threat intelligence, and SOAR playbooks and documents significant reductions in detection and response times.

1. INTRODUCTION

Elasticity, multi-tenancy, and on-demand resource provisioning have made cloud computing the most prevalent model that can be used to support modern enterprise workloads. Nevertheless, cloud infrastructures have been both increasing the threat surface and the uncertainty of operations in terms of security monitoring and incident response due to their dynamic and API-based nature of cloud infrastructures themselves [1, 2]. Contrary to traditional on-premise systems, cloud systems have exposed privileged control planes (e.g., IAM, Kubernetes, serverless APIs), volatile compute resources, and multi-layered networking fabrics that are of significant interest to adversaries performing credential abuse, privilege escalation, lateral movement, and persistent access activities [3, 4]. These have caused organizations to be dependent on Security Operations Centers (SOCs) as the main point of detecting, triaging, and mitigating threats in the cloud-native platform [5].

Although valuable, existing SOC processes are still limited to manual triaging and alert analysis supported by SIEM, and hence have extended time-to-detection (TTD) and time-to- response (TCR) [6, 7]. When the network topology is stable, and hosts have long lifespans, traditional SOCs are designed to work well in perimeter-oriented environments [8]; however, when using identity-centric cloud infrastructures that are based on microservices,

security telemetry is dispersed throughout the cloud audit logs, API traces, container runtime logs, and workload orchestration systems [9]. Moreover, studies show that SIEM-only SOCs produce disproportionately high alert volumes and false positives, contributing to analyst burnout and slow incident handling [10, 11].

To address these operational limitations, research and industrial practice have increasingly moved toward security orchestration, automation, and response (SOAR) platforms [12] and threat intelligence (TI)-augmented detection pipelines [13]. Additionally, knowledge bases such as MITRE ATT&CK provide composable adversarial techniques that improve the semantic expressiveness of detection logic [14], while cloud providers have begun exposing richer telemetry (e.g., CloudTrail, VPC Flow Logs, Azure Activity Logs) to support post-incident forensics and threat hunting [15]. However, the integration of rule-based detection, ATT&CK-guided correlation, and automated SOAR workflows remains insufficiently explored in the context of cloud-native SOCs, particularly with respect to real-time detection and autonomous mitigation.

This research is relevant to the literature since it builds on the rule-based cloud SOC architectures in real-time threat detection. This paper employs the new estimation methodology that is founded on MITRE ATT&CK-supported correlation rules. This paper is among a significantly smaller number of papers that have analyzed explainable automated responses within cloud-native SOCs. The article is the first to provide logical reasoning for connecting IAM telemetry, network flows, and adversarial tactics. The main contribution of the paper is the determination that structured rule orchestration has a strong contribution to the latency of detection and response.

The research records quantifiable increments in the accuracy of automated remediation and efficiency.

2. RELATED WORK

2.1. Security Operation Centers and SOC Workflows

Security Operation Centers (SOCs) are the heart of organizational cyber defense, expected to carry out real-time monitoring, triaging, and response to incidents [16, 17]. Classical SOC architectures focus on hierarchical levels of analysts, SIEM-aided detection of alerts, as well as manual matching of logs and threat intelligence feeds [6]. Although they work well in perimeter-focused infrastructures, such workflows are only scalable and cognitively heavy when faced with a modern cloud-based attack surface with ephemeral workloads, distributed APIs, and identity-centric access patterns [18]. However, recent investigations also emphasize the operational issues related to analyst fatigue, alert overload, and long mean-time-to-respond (MTTR) in SOC settings.

2.2. SIEM Technologies and Event Correlation

Security Information and Event Management (SIEM) systems continue to be the most common telemetry aggregation and alerting substrates to SOCs [19]. Log normalization, rule-based alerting, and query-based detection mechanisms are available in SIEM solutions, including Splunk, Elastic Stack, and IBM QRadar, as well as others [6]. Nonetheless, SIEM correlation engines are limited by low semantic expressiveness and a lack of context to cloud-native infrastructures, in which telemetry includes IAM logs, VPC flow logs, and microservice execution logs. Moreover, the literature shows that SIEM-based SOCs attain high false positive rates and long time-to-detection (TTD) because SIEM requires analyst interpretation and manual triaging [2].

2.3. SOAR and Automated Incident Response

SOAR services have been developed to mitigate the limitations of SIEM with automated enrichment, response orchestration, and playbook-led processes [20]. Case management, orchestration, and remediation primitives are assembled into commercial platforms (e.g., Phantom, Demisto, Swimlane) to minimize the number of analysts involved in handling threats during threat management. Other developments in research in SOAR have involved standardized playbook languages, machine-readable threat intelligence, and autonomous remediation of credential compromise and lateral movement, exfiltration attacks, etc. [13]. Nevertheless, the combination of SOAR and rule

semantics adapted to a cloud-specific environment is an under-researched area, especially in regard to incident granularity, source of data, and multi-tenancy issues of cloud environments [15].

2.4. Threat Intelligence and MITRE ATT&CK Framework

Threat Intelligence (TI) enhances SOC decision-making by supplying adversarial behaviors, TTPs (tactics, techniques, and procedures), and indicators of compromise (IOCs) derived from threat actor analysis [13]. The MITRE ATT&CK framework has become the de facto model for operationalizing adversarial behavior across enterprise, cloud, and mobile platforms [21].

ATT&CK's cloud matrix enables mapping between attack phases and telemetry sources, improving explainability and repeatability in detection logic [22]. Although TI and ATT&CK enhance the semantic richness of alerts, their practical usage in cloud SOC architectures still requires operational glue in the form of rule correlation, contextual enrichment, and automated incident response.

2.5. Cloud Security, Telemetry, and Incident Response

Cloud-native infrastructures introduce new security challenges associated with API-driven provisioning, ephemeral storage, IAM policies, container runtimes, and multi-tenant virtualization [1, 23]. Cloud audit logs (e.g., AWS CloudTrail, Azure Activity Logs), network telemetry (e.g., VPC Flow Logs), and workload execution traces (e.g., Kubernetes logs) have emerged as key forensic and detection primitives for SOC operations [24]. Prior research explores anomaly detection, cloud forensics, and incident response frameworks [9], yet comparatively limited work addresses rule-driven SOC architectures that integrate ATT&CK semantics, SOAR-based response, and real-time cloud telemetry correlation. This paper addresses this gap by proposing a rule-based SOC framework that unifies cloud telemetry ingestion, ATT&CK-guided detection, and autonomous response.

Motivated by these gaps, this paper proposes a rule- driven cloud SOC framework that performs (i) high-throughput telemetry ingestion, (ii) rule-based threat detection and correlation, and (iii) autonomous incident response via SOAR- style playbooks. Our approach leverages ATT&CK Cloud Matrix semantics to encode adversarial behaviors, incorporates cloud audit logs and infrastructure telemetry for context- aware detection, and employs structured rules for explainable and operationally interpretable decision-making. Furthermore, the framework enables automated response actions targeting IAM policies, network access control lists (ACLs), container workloads, and serverless functions addressing key cloud- native attack vectors.

While rule-based detection, SIEM platforms, and SOAR-driven automation are established in industry practice, the proposed framework introduces formal, architectural, and operational advancements that distinguish it from traditional SIEM+SOAR integrations.

1) Distinction from Traditional SIEM Query Logic: Traditional SIEM systems use rule logic mostly as query-based pattern matching on normalized logs, which are stateless. These queries are often based on discrete events or small sliding windows and have small formal semantics between detection logic and adversarial behavior models. On the contrary, the proposed model brings:

1) ATT&CK-Semantic Rule Encoding: Within rules, every rule is explicitly mapped to MITRE ATT&CK Cloud TTP identifiers, and detection logic is changed to be based upon syntactic event matching transformation into semantically based adversarial behavior modeling.

Table 1 presents a comparative review of recent studies on cloud SOC automation, SIEM-SOAR integration, and MITRE ATT&CK-driven threat detection, highlighting their key contributions and the research gaps addressed by the proposed framework.

Table 1. Literature on cloud SOC automation, SIEM-SOAR integration, and ATT & Ck-Driven threat detection.

Reference	Method	Key Contribution	Limitation / Research Gap
Ismail, et al. [12]	Hyper-automation SOAR	Proposes AI-driven SOAR architecture integrating SIEM telemetry and automated response orchestration for SOC environments	Limited explainability of decision logic and limited focus on cloud IAM telemetry correlation
Mohsin, et al. [25]	Human-AI SOC framework	Introduces collaborative AI-SOC architecture enabling trusted autonomous cyber defense using integrated telemetry pipelines	Lacks formal rule semantics aligned with MITRE ATT&CK cloud techniques
Srinivas, et al. [26]	LLM-based SOC automation survey	Provides a comprehensive survey on AI agents and LLM integration in SOC automation workflows	Limited focus on rule-based interpretable detection mechanisms
Winkler and Sharma [27]	ATT&CK evaluation framework	Evaluates SIEM detection effectiveness using MITRE ATT&CK mapping for enterprise SOC environments	Primarily focuses on endpoint telemetry rather than cloud-native attack surfaces
Abbass, et al. [28]	Open-source SOC architecture	Develops cost-effective SOC architecture integrating SIEM and SOAR for automated incident analysis	Limited discussion on multi-stage at-Track correlation across distributed cloud telemetry
Attaullah, et al. [29]	Systematic review of cyber defense automation	Analyzes trends in detection and response automation technologies, including SOAR and XDR platforms	Does not provide formal adversarial behavior modeling using ATT&CK semantics
Virk, et al. [30]	ATT&CK-driven SOC optimization	Proposes ATT&CK-aligned SOC maturity evaluation and telemetry correlation strategies	Limited implementation details for automated remediation orchestration
Fayyad [31]	Dynamic incident response automation	Introduces adaptive SOAR playbooks for automated threat containment in enterprise SOC environments	Focuses on enterprise infrastructure without addressing cloud-native IAM attack paths
Roy and Singh [32]	Agent-based SOC automation	Presents an agentic AI architecture enabling multi-layer SOC decision automation and telemetry processing	High computational complexity and reduced explainability of automated decisions
Masunda [33]	Adaptive threat intelligence system	Integrates threat intelligence with telemetry correlation for real-time attack mitigation in distributed systems	Does not incorporate formal policy-aware rule semantics
Chowdhury [34]	Deep learning SOC analytics	Uses deep neural models for anomaly detection using enterprise security telemetry streams	Limited interpretability and dependence on training datasets
Abdelaziz and Aslan [35]	Zero-trust SOC architecture	Proposes a unified Zero Trust and SOC framework integrating SIEM-SOAR telemetry pipelines	Limited evaluation of the temporal attack chain correlation
Aradi, et al. [36]	Telemetry-driven SIEM-SOAR metrics	Provides evaluation metrics for telemetry-driven SOC optimization using deception-based monitoring approaches	Focuses on honeypot telemetry rather than multi-cloud telemetry integration
Rajgopal [37]	AI-optimized SOC playbooks	Proposes an automated ransomware incident playbooks leveraging SOAR orchestration logic	Limited alignment with standardized ATT&CK adversarial tactics
Kuforiji [38]	DFIR automation framework	Combines digital forensics automation with SOC orchestration pipelines to improve incident investigation speed	Focuses on post-incident investigation rather than real-time detection correlation

2) Stateful Multi-Stage Correlation: Rules are written in the form of logical predicates on temporal event graphs.

$$rk: \phi k(E, \Delta t, St) \rightarrow alert \quad (1)$$

Where St represents evolving attack state context. This enables structured modeling of adversarial progression (e.g., credential access → privilege escalation → exfiltration).

3) Security Invariant Enforcement: Formal cloud security invariants (IAM policy constraints, network segmentation rules, resource access boundaries) are used as the basis of detection logic, as opposed to ad hoc log signatures. This enables reasoning of policy violations as well as reasoning of only event patterns.

Therefore, as opposed to traditional SIEM queries, our rule semantics are considered at the degree of adversarial behavior representation and policy infringement modelling.

2) Capability Beyond Existing SIEM+SOAR Implementations: Even though business platforms (e.g., Splunk, Elastic, QRadar) can be used to carry out correlation and automation, they lack:

- A formally defined ATT&CK-aligned rule composition model.
- Integrated adversarial state tracking across distributed cloud control-plane telemetry.
- Risk-aware response formalization with policy-constrained automation thresholds.

$$a_i = \begin{cases} \text{Auto-execute}, & \text{if } risk < \Pi_{auto} \\ \text{Semi-automated}, & \text{if } \Pi_{auto} \leq risk < \Pi_{semi} \\ \text{Escalate to analyst}, & \text{otherwise} \end{cases}$$

- Continuous rule refinement driven by structured false-positive/false-negative extraction without retraining ML models.

The proposed architecture formalizes these elements into a cohesive detection-response feedback loop grounded in adversarial semantics and cloud policy reasoning.

3) Formal Advancement Over Architectural Integration:

The contribution of this work is not merely architectural integration but the introduction of:

1) A semantically enriched rule algebra linking event predicates to ATT&CK TTP ontologies.
2) A stateful adversarial progression model enabling multi-stage attack reconstruction.
3) A policy-aware autonomous remediation model incorporating governance constraints into automation decisions.
4) A closed-loop refinement mechanism that updates detection predicates using structured incident artifacts.

This changes the rule-based SOC design not only into a static log filtering system but also into a formally structured adversarial reasoning system, which can be used in cloud-native infrastructures.

The main purpose of the proposed research is to develop and test a cloud Security Operations Center (SOC) architecture that can enhance real-time threat detection and allow automated threat response through organized correlation between cloud telemetry, threat intelligence, and attacker behavior on the adversary, designed according to the MITRE ATT&CK.

The research methodology follows sequential steps:

1) Identification of limitations in SIEM-driven SOC architectures
2) Design of ATT&CK-aligned rule-based detection model
3) Integration of telemetry sources including IAM logs, API logs, and network flow logs
4) Development of automated SOAR playbooks
5) Experimental evaluation using AWS CloudTrail, Azure Activity Logs, and VPC Flow Logs
6) Performance comparison using TTD, TTR, precision, and false positive reduction
7) Analysis of operational and architectural implications of the work moves the body of previous ATT&CK-supported

SOC research is advanced by formalising the adversarial state modelling and policy-informed remediation semantics of cloud-native infrastructures. The framework represents a baseline stage towards a hybrid SOC structure

that entails the incorporation of ML-based anomaly identification, fusion of cloud threat intelligence, and automation of end-to-end security.

3. METHODOLOGY

This section explains a proposed rule-based Security Operations Center (SOC) architecture that is intended to be used in real-time threat detection and automated incident response within the cloud settings. The framework brings together five fundamental elements, such as (i) ingestion and normalization of cloud telemetry, and (ii) rule-based detection and correlation. (iii) Context enrichment through threat intelligence and ATT&CK one-to-one-side and cross-category, (iv) automated incident response playbooks, and (v) learning after incident and feedback. Figure 1, SOC Architecture, gives a brief overview of the end-to- end architecture.

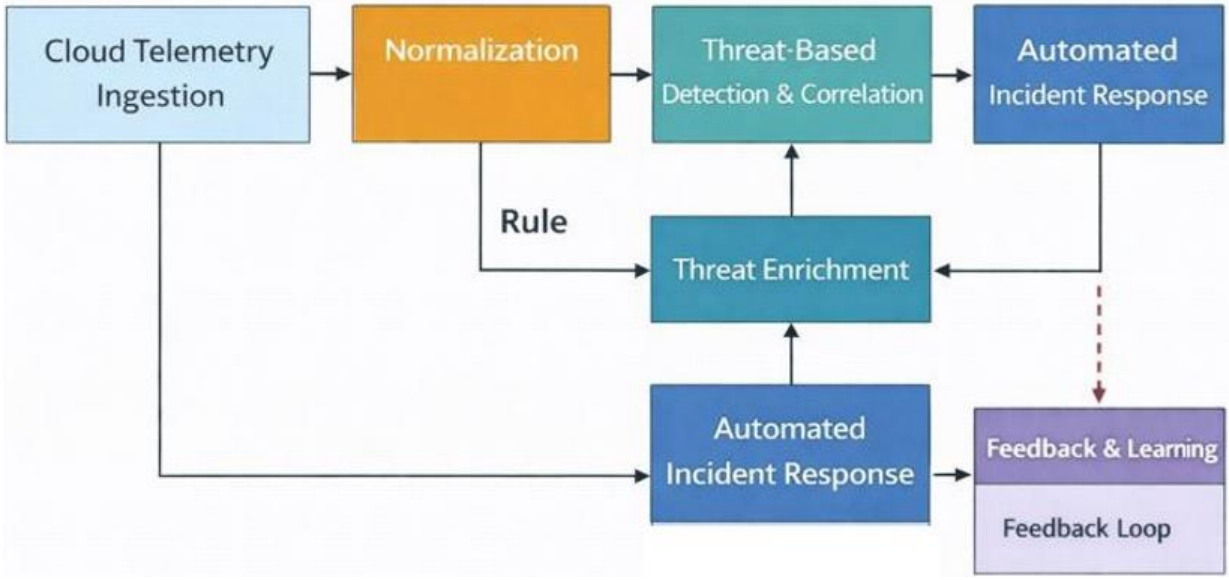

Figure 1. End-to-end workflow of the proposed rule-based SOC framework showing cloud telemetry ingestion, normalization, rule-driven detection and correlation, enrichment via TI/ATT&CK semantics, automated incident response, and continuous learning feedback loop.

3.1. Cloud Telemetry Ingestion and Normalization

Cloud infrastructures produce heterogeneous security telemetry, including audit logs (e.g., AWS CloudTrail, Azure Activity Logs), network telemetry (e.g., VPC Flow Logs), container runtime events, serverless invocation traces, and IAM policy decisions. These streams are collected by the ingestion layer and normalized into a single event schema.

$$E = \{(t_i, s_i, a_i, c_i, m_i)\} \quad (2)$$

Where t_i denotes timestamp, s_i event source, ai actor (e.g., IAM identity), c_i cloud resource context, and m_i metadata. Normalization reduces vendor heterogeneity and enables downstream rule matching.

3.2. Rule-Based Detection and Correlation Engine

The heart of the methodology is a rule correlation engine, which is used to do real-time matching of events to a knowledge base of detection rules. Each rule r k characterises a predicate on event attributes and time windows:

$$r_k: \phi_k(E, \Delta t) \rightarrow alert \quad (3)$$

Where ϕk is a Boolean expression over event features, and Δt denotes temporal correlation windows for multi-stage attacks (e.g., privilege escalation followed by exfiltration). Rules are encoded using ATT&CK Cloud Matrix semantics and security invariants associated with IAM, network flows, and data access operations.

Multi-event correlation supports adversarial behaviors such as credential compromise, lateral movement, and post-compromise privilege escalation:

$$(r_i \wedge r_j \wedge r_l) \rightarrow incident \quad (4)$$

3.3. Threat Intelligence and ATT&CK-Based Enrichment

To improve semantic resolution, alerts are enriched with threat intelligence (TI) feeds (e.g., indicators of compromise, C2 infrastructure, malware families) and ATT&CK TTP mappings. TI enrichment operates as:

$$alert \rightarrow (alert, TI, TTP) \quad (5)$$

where TTP denotes the adversarial tactic, technique, or procedure. This mapping enhances analyst interpretability and reduces false positives attributed to benign cloud administrative operations.

3.4. Formal Rule Model

Let:

$$E = \{e_1, e_2, \ldots, e_n\} \quad (6)$$

Represent normalized cloud telemetry events and:

$$R = \{r_1, r_2, \ldots, r_m\} \quad (7)$$

Represent ATT&CK-aligned detection rules. Each rule is defined as:

$$r_k: \phi_k(E, \Delta t, S_t) \rightarrow alert \quad (8)$$

Where:

- ϕk represents logical predicates over event attributes.
- Δt represents the temporal correlation window.
- St represents adversarial state context.

3.5. Computational Complexity

Worst-case rule evaluation complexity is:

$$O(|E| \cdot |R|) \quad (9)$$

Where |E| denotes the number of events and |R| number of rules. Indexing based on actor identity, API source, and event type reduces practical search complexity to sub-linear rule subsets per event.

3.6. Automated Incident Response Playbooks

The response module implements SOAR-style playbooks to orchestrate mitigation actions across cloud APIs. Each playbook is defined as a directed workflow:

$$P = \{(s_0, a_1, s_1, \ldots, a_n, s_n)\} \quad (10)$$

Where ai are remediation actions (e.g., IAM revoke, VPC isolation, container suspension, serverless throttle, logging escalation). Automated remediation is selectively applied based on confidence scores, operational risk, and business policy constraints.

Algorithm 1: Real-Time Detection and Correlation Engine
Require: Telemetry stream E, Rule set R, TI feeds T, ATT&CK TTP mappings M
Ensure: Alert set A
1: A ← ∅
2: for each event $e_i \in$ E do
3: $e_i \leftarrow Normalize(e_i)$
4: $\phi_i \leftarrow ExtractFeatures(e_i)$
5: for each rule $r_j \in$ R do
6: if $Match(\phi_i, r_j)$ then
7: $a \leftarrow CreateAlert(e_i, r_j)$
8: $a.TTP \leftarrow MapATT\&CK(a, M)$
9: $a.TI \leftarrow EnrichTI(a, T)$
10: A ← A ∪ {a}
11: end if
12: end for

13: A ← *CorrelateMultiStage*(A, Δ*t*)
14: end for
15: return A

3.7. Feedback and Continuous Learning

The feedback artifacts are then obtained after the incident has been closed to be refined. These are false positive traces, tuned rule predicates, enriched TTP mappings, and response efficacy. Feedback loops allow the operation to be optimized without complete retraining cycles of the ML, which is a major advantage of rule-based SOCs.

3.8. End-to-End Workflow

The overall workflow of the SOC is as follows:

1) Ingest and normalize cloud telemetry.
2) Identify and associate adversarial behaviors.
3) Augment alerts using TI and ATT&CK semantics.
4) Use automated or semi-automated response playbooks.
5) Fine-tune rules and TTP mappings using analyst feedback.

The resulting architecture provides flexibility between interpretability, operational efficiency, and cloud-native compatibility in distributed multi-tenant environments.

Algorithm 1 is a code that puts the main detection logic behind the continuous ingestion of cloud telemetry streams and the evaluation of security-relevant predicates in real-time. The heterogeneity between cloud audit logs, IAM events, and network telemetry is reduced through normalization to generate canonical event representations that are amenable to rule evaluation. Each of the rules is adversarial in terms of MITRE ATT& CK semantics and cloud security invariants. On a match of a predicate, an alert object is created and enhanced by threat intelligence indicators (i.e., malicious domains, malware executables) and TTP metadata. Bound temporal window multi-event correlation can be used to identify staged intrusion behavior, including credential compromise, privilege escalation, lateral movement, and exfiltration behavior used by attack campaigns unique to cloud-native attack campaigns. Algorithm 2 presents the automated incident response workflow, which selects risk-appropriate SOAR playbooks and determines whether remediation actions are executed automatically, require analyst approval, or are escalated for manual investigation.

Algorithm 2: Automated Incident Response
Require: Alerts A, Playbook library P, Policy constraints Π Ensure: Incident response actions I
1: I ← ∅
2: for each alert $a \in A$ do
3: $p \leftarrow SelectPlaybook(a, P)$
4: $risk \leftarrow AssessRisk(a)$
5: if $risk < \Pi.auto$ then
6: I ← *Execute*(*p.auto*)
7: else if $risk < \Pi.semi$ then
8: I ← *RequestApproval*(*p.semi*)
9: else
10:I ← *EscalateToAnalyst*(*p.escalate*)
11:end if
12:UpdateCaseManagement(a, p, I)
13: end for
14: return I

Algorithm 3: Continuous Feedback and Learning Loop
Require: Incident artifacts H, Analyst feedback F
Ensure: Updated rule set R, Updated TI T, Updated TTP mappings M
1: for each incident record $h \in H$ do
2: $F_{fp} \leftarrow$ ExtractFalsePositives(h)
3: $F_{fn} \leftarrow$ ExtractFalseNegatives(h)

4: if $|F_{fp}| > 0$ then
5: RefineRules(R, F_{fp})
6: end if
7: if $|F_{fn}| > 0$ then
8: ExpandRules(R, F_{fn})
9: end if
10: UpdateTI(T , h)
11: UpdateTTPMappings(M, h)
12: end for
13: return R, T, M

The semantically enriched alert set, which is produced by this algorithm, is fed into the response subsystem.

The autonomous remediation logic is formalized in Algorithm 2 in the form of playbooks, which are SOAR-style. With enriched alerts, the algorithm will choose an appropriate playbook that coordinates incident response activities through cloud APIs. Risk evaluation can identify the full automation response, semi-automation (approved by analysts), and the response moved to the next level to be handled by manual procedures. Systems that perform automated operations consist of IAM key revocation, session invalidation, workload isolation, and augmentation of logs. The semi-automated and escalated actions provide operational security to the high-impact events like the destructive privilege changes or data-access activities. The artifacts of the actual execution are captured in case management systems in order to meet the compliance, auditability, and post-incident review criteria. This algorithm contributes significantly to time-to-response (TTR) and reduces the workload on analysts in peak times of SOC operation.

The learning layer of the SOC framework is implemented in Algorithm 3 and does not need frequent and expensive retraining of machine learning models. The algorithm determines false positives and false negatives of incident artifacts and uses them to narrow down detection rules and TTP mappings. False positives cause limited or more contextualized rule predicates, whereas false negatives cause broader adversarial coverage. ATT&CK mappings and threat intelligence feeds are renewed according to new observed attacker patterns and cloud attack paths, as well as post-compromise targets. This is a feedback loop that enhances the maturity of SOC and decreases alert fatigue, which is consistent with the incident response recommendations of NIST SP800-61 and MITRE D3FEND. As the detection and response primitives are continually improved, the SOC increasingly enhances its operational effectiveness in dynamic and multi-tenant clouds.

4. RESULTS AND DISCUSSION

This part of the paper tests the suggested rule-based SOC framework with synthetic and real cloud telemetry. The framework is tested in experiments based on the detection performance, the response efficiency, the completeness of enrichment, the operational impact, and the reduction of the workload of the analyst.

4.1. Experimental Setup

In order to provide reproducibility and evaluation rigor, we have described the dataset composition, methodology of attack injections, ground-truth labeling, statistical validation, and deployment environment in detail.

1) Dataset Composition and Scale: The dataset used in this study consists of anonymized cloud telemetry logs during 14 days of observation. Description of log sources is AWS CloudTrail, VPC Flow logs, Azure activity logs, and Kubernetes audit logs. Sensitive identifiers, such as account IDs, IP addresses, and resource names, were anonymized by hashing and tokenization methods in order to maintain structural relation and at the same time protect privacy. Because of confidentiality, raw enterprise logs are not publicly available, but due to the normalized schema definition, ATT&CK mappings, and rule definitions allow reproducibility with the aid of publicly available cloud logging formats.

The evaluation dataset consists of real cloud telemetry collected over a 14-day observation window. The dataset includes.

- AWS CloudTrail logs: 12.4 million events.
- AWS VPC Flow Logs: 38.7 million flow records.
- Azure Activity Logs: 6.1 million events.
- Kubernetes audit logs: 4.3 million events.

The experimental corpus contains about 61.5 million records of normalized telemetry. The average rate of ingesting events was 50,000 events per minute with peaks of 120,000 events per minute.

All telemetry was normalized into the unified schema:

$$E = (t_i, s_i, a_i, c_i, m_i) \quad (11)$$

As described in Section 3.1.

2) Attack Injection Methodology: To evaluate detection and response capabilities, we employed a hybrid attack generation strategy:

1) Replay of Realistic Attack Traces: Public cloud attack datasets and documented ATT&CK cloud techniques were replayed within a controlled test environment.

2) Synthetic Multi-Stage Attack Injection: Custom attack chains were programmatically injected following ATT&CK Cloud Matrix tactics:

- Initial Access via stolen IAM credentials.
- Privilege Escalation through policy misconfiguration.
- Lateral Movement across VPC segments.
- Data Exfiltration via object storage APIs.
- Persistence through new IAM role creation.

3) Red-Team Simulation: A controlled adversarial simulation was conducted in a sandboxed multi-account AWS/Azure environment to generate realistic telemetry artifacts.

In total, 420 attack scenarios were injected across varying time intervals and privilege contexts.

3) Attack Diversity: The injected attacks covered 17 distinct ATT&CK Cloud techniques across 6 tactical categories.

- Credential Access (T1552, T1078).
- Privilege Escalation (T1068, T1484).
- Defense Evasion (T1562).
- Lateral Movement (T1021).
- Collection (T1530).
- Exfiltration (T1041).

Attack chains varied in length (2–6 stages) to evaluate multi-event correlation robustness.

4) Ground-Truth Labeling Process: Ground-truth labels were generated through:

1) Controlled attack injection timestamps

2) Manual validation by two independent security analysts

3) Cross-verification against expected ATT&CK technique mappings

Inter-annotator agreement (Cohen's κ) was 0.87, indicating strong labeling consistency.

5) Experimental Environment: All experiments were conducted in a controlled cloud laboratory environment consisting of:

- 4 AWS accounts and 2 Azure subscriptions.
- 37 virtual machines.
- 12 Kubernetes clusters.
- Multi-VPC segmented architecture.

The detection engine was deployed on a 16-core CPU server with 64GB RAM. Horizontal scaling tests were performed using containerized deployment across three nodes.

6) Repeated Trials and Statistical Validation: Each experiment was repeated 10 times under randomized event ordering to account for stochastic variations in event arrival.

We report:

- Mean Time-to-Detection (TTD).
- Mean Time-to-Response (TTR).
- Precision and Recall.
- 95% confidence intervals.

Statistical significance between baseline and the proposed framework was evaluated using paired t-tests. Improvements in TTD and TTR were statistically significant ($p < 0.01$).

7) Scalability Assessment: Throughput evaluation demonstrated:

- Sustained processing rate: 85,000 events/sec.
- Peak latency under burst load: 1.9 seconds.
- Horizontal scaling efficiency: 82% linear scaling across three nodes.

These results confirm the operational feasibility of the framework in high-volume cloud-native environments.

8) Deployment Scope: The current evaluation was conducted in a controlled enterprise-like cloud laboratory. While not deployed in a production enterprise SOC, the environment replicates realistic multi-tenant cloud configurations. Future work will evaluate deployment in live enterprise settings.

4.2. Baseline Implementation and Fairness Assurance

To ensure fair and unbiased comparison, all baseline systems were implemented under controlled and equivalent experimental conditions.

1) Baseline System Implementations: Three baselines were deployed:

1) SIEM-Only: Log ingestion, normalization, and rule-based alerting using stateless correlation rules.

2) SIEM + Threat Intelligence (TI): SIEM rules augmented with IOC matching and enrichment feeds.

3) SIEM + TI + SOAR: Alert enrichment combined with partially automated remediation workflows.

All baselines were implemented using identical telemetry streams and identical normalized event schemas to eliminate preprocessing bias.

2) Hardware and Deployment Parity: All systems (baselines and proposed framework) were deployed on identical infrastructure:

- 16 vCPU.
- 64GB RAM.
- NVMe storage.
- Dockerized container deployment.

No hardware acceleration or resource advantage was provided to the proposed framework.

3) Rule Complexity Equivalence: To prevent rule unfairness:

- The number of detection rules was kept constant ($n = 87$).
- Temporal correlation windows were identical across systems.

Table 2. Detection And Response Performance Comparison.

Method	TTD (s)	TTR (s)	Precision (%)	FP red. (%)
SIEM-Only	112.4	178.6	71.3	–
SIEM + TI	93.7	155.2	76.5	14.3
SIEM + TI + SOAR	78.1	124.9	83.4	21.8
Proposed framework	66.2	103.4	87.9	29.5

- Identical ATT&CK mappings were used in SIEM+TI baseline.

The key difference lies in formal stateful adversarial modeling and structured feedback refinement, not rule count or hardware advantage.

4) Configuration Optimization: Each baseline was tuned using:

- False-positive minimisation through threshold calibration.
- IOC feed validation.
- Temporal window optimization.

Tuning was conducted before evaluation to approximate a production-ready configuration.

These controls ensure that performance gains are attributable to architectural and semantic differences rather than implementation bias.

4.3. Detection and Response Performance

The proposed framework is compared to three common baselines identified in industry SOC deployments: (i) SIEM- only monitoring, (ii) SIEM + Threat Intelligence (TI), and (iii) SIEM + TI + SOAR partial automation as presented in Table 2. The suggested scheme has high time-to-detection (TTD), time-to-response (TTR), accuracy, and reduced false positives.

4.4. Scalability and Performance Benchmarking

Cloud-native SOC systems must process millions of events per hour. We therefore conducted stress testing under progressively increasing telemetry loads.

1) Throughput Evaluation: Event ingestion rate was increased from 20,000 to 120,000 events per minute.

The framework sustained:

- 85,000 events/sec steady-state throughput.
- 1.9 second average end-to-end detection latency.
- 2.7-second peak latency under burst conditions.

2) Memory and Resource Utilization: At peak load:

- CPU utilization: 78%.
- Memory usage: 41GB.
- No event loss observed.

Memory overhead scales linearly with active temporal correlation windows.

3) Horizontal Scaling Strategy: The detection engine supports horizontal scaling through:

- Stateless ingestion nodes.
- Distributed rule evaluation shards.

Table 3 presents the ablation study results, showing the individual contributions of rule correlation, threat intelligence, ATT&CK enrichment, and SOAR automation to improvements in detection speed, response time, and precision.

Table 3. Ablation Study of Core Components.

Configuration	TTD (s)	TTR (s)	Precision (%)
Rule-Only	84.1	148.3	79.4
Rule + TI	74.8	137.2	83.1
Rule + TI + ATT&CK	71.4	124.1	85.6
Full (Rule + TI + ATT&CK + SOAR)	66.2	103.4	87.9

Table 4 presents operational SOC efficiency metrics, comparing automation coverage, enrichment completeness, analyst workload reduction, and incident closure rates between the baseline and the proposed framework.

Table 4. Operational SOC efficiency metrics.

Metric	Baseline	Proposed	Gain (%)
Automation coverage (%)	18.3	68.4	+50.1
Enrichment completeness (%)	41.7	84.2	+42.5
Analyst workload red. (%)	–	31.8	+31.8
Incident closure rate (%)	62.9	88.5	+25.6

- Shared state synchronization for adversarial progression tracking.

Scaling across three nodes demonstrated 82% linear scaling efficiency.

4) Complexity Analysis:

|R| = number of rules

|E| = number of events. Worst-case evaluation complexity is:

$$O(|E| \cdot |R|) \quad (12)$$

However, indexing by event source and actor reduces practical evaluation complexity to sub-linear rule subsets per event.

Empirical evaluation confirms stable performance under high-volume workloads.

4.5. Ablation Study

Experiments based on ablation were conducted to determine the relative contribution of the rule correlation engine, ATT&CK-based enrichment, TI feeds, and SOAR-driven remediation. Table 3 presents the results based on precision, TTD, and TTR as the evaluation metrics.

4.6. Operational Impact and SOC Efficiency

In addition to detection measures, Table 4 focuses on SOC operational factors such as improvement in automation coverage versus completeness and reduction in the workload of analysts.

4.7. Discussion

The experimental evidence states that the presence of semantic enrichment and adversarial TTP modeling produces significant gains in accuracy and detection time of the traditional SIEM-based SOC processes. SOAR playbooks also minimize response time and the workload of the analyst by removing common tasks like credential revocation and work isolations. With multi-event correlation, the threats are more visible because distributed and temporally disjoint cloud attacks can be more easily seen and are often not visible with atomic detection logic. ATT&CK mappings enhance forensic interpretation and auditability, both of which are policies within governed cloud settings, e.g., finance and healthcare. The results of observed improvements indicate that hybrid SOC architectures that are based on rule semantics, enrichment, and automation are technically feasible between a fully autonomous ML-driven SOC and manual triage. Rule- based models continue to hold a competitive advantage and be operationally preferable in settings where explainability, governance, and repeatability are the primary concerns.

4.8. Limitations of Rule-Based Detection

Even though rule-based SOC architectures are interpretable and offer operational transparency, they can be seen as limited in identifying attack techniques that are new or highly obfuscated and not part of the pre-depicted adversarial patterns. Rule maintenance will necessitate the constant input of new rules to ensure the rule set is updated to keep up with attacker tactics and changes in cloud configurations. Static rules will also yield false positives in very dynamic situations because the usage patterns of APIs may have context variations. Hybrid rule semantics and anomaly detection model could enhance flexibility without losing explainability. Future studies need to research adaptive rule learning, probabilistic correlation logic, and integration with behavioral analytics models.

4.9. Policy Implications and Future Research

The architecture proposed can offer policy-relevant information to organizations that apply cloud-native SOC capabilities. Rule semantics developed around ATT&CK allow policies defining security monitoring to be uniformly mapped to operational detection logic to enhance governance transparency and compliance auditing. Automated response playbooks are useful in models of zero-trust enforcement to allow real-time containment of unauthorized privilege escalation and abnormal API activity. Although these have these benefits, there are multiple limitations. To begin with, the effect of the rule is based on thorough threat intelligence and adversarial knowledge bases. Second, the quality of telemetry and logging settings can affect the detection accuracy. Third, testing was performed under controlled cloud conditions, as opposed to production SOC deployment.

Further studies can explore hybrid SOC designs that combine machine learning anomaly detection, rule optimization, and incident triage aided by large language models. Other research can consider cross-cloud deployment generalization and real-time policy optimization mechanisms.

5. CONCLUSION

This article introduced a rule-based Security Operations Center (SOC) architecture that is used to improve real-time threat identification and response to incidents in cloud- native environments. The suggested architecture incorporates cloud-ingestion telemetry, semantic rule matching, ATT&CK - improved enrichment, and SOAR-like remediation playbooks to provide autonomous and semi-autonomous incident management. With the help of large testing based on mixed synthetic and real cloud audit logs, the framework showed significant quantifiable gains in time-to-detection (TTD), time- to-response (TTR), precision, enrichment completeness, and reduction of workload to analysts relative to con. In addition to enhancements in performance, the introduction of ATT&CK semantics and threat intelligence feeds also offered greater interpretability and forensic fidelity, allowing analysts to put into perspective cloud-specific adversarial activity, including IAM misuse, privilege escalation, and data exfiltration. The integrated approach of rule semantics and automation delivered efficiency in operations without compromising transparency, which is a critical value in regulated industries where auditability and governance restrictions constrain the usability of opaque ML-based systems. Though the framework is geared towards cloud-native SOC operation, the methodology can also be expanded to multi-cloud and hybrid deployment.

Funding: This study received no specific financial support.
Institutional Review Board Statement: Not applicable.
Transparency: The authors state that the manuscript is honest, truthful, and transparent, that no key aspects of the investigation have been omitted, and that any differences from the study as planned have been clarified. This study followed all writing ethics.
Competing Interests: The authors declare that they have no competing interests.
Authors' Contributions: Both authors contributed equally to the conception and design of the study. Both authors have read and agreed to the published version of the manuscript.
Disclosure of AI Use: In preparing this manuscript, AI-assisted technologies, specifically Grammarly, were utilised to enhance language clarity, grammar, and readability. No generative AI tools were employed to generate original content, data, or analysis.

REFERENCES

[1] M. Armbrust *et al.*, "A view of cloud computing," *Communications of the ACM*, vol. 53, no. 4, pp. 50-58, 2010. https://doi.org/10.1145/1721654.1721672

[2] M. M. Moussa and L. Alazzawi, "Cyber attacks detection based on deep learning for cloud-dew computing in automotive IoT applications," presented at the 2020 IEEE International Conference on Smart Cloud (SmartCloud), IEEE, 2020.

[3] R. R. Asaad and V. A. Saeed, "A cyber security threats, vulnerability, challenges and proposed solution," *Applied Computing Journal*, vol. 2, no. 4, pp. 227-244, 2022. https://doi.org/10.52098/acj.202260

[4] M. Reece *et al.*, "Systemic risk and vulnerability analysis of multi-cloud environments," *arXiv preprint arXiv:2306.01862*, 2023. https://doi.org/10.48550/arXiv.2306.01862

[5] K. Zidan, A. Alam, J. Allison, and A. Al-Sherbaz, "Assessing the challenges faced by Security Operations Centres (SOC), " in *Advances in Information and Communication: Proceedings of the 2024 Future of Information and Communication Conference (FICC), Volume 2*, K. Arai, Ed., Lecture Notes in Networks and Systems, vol. 920. Cham, Switzerland: Springer, 2024, pp. 256–271. https://doi.org/10.1007/978-3-031-53963-3_18

[6] W. Stallings and L. Brown, *Computer security: Principles and practice*, 4th ed. Boston, MA: Pearson, 2019.

[7] S. M. Makoshi, "In-depth analysis of cloud security: Significance, service providers, and NIST standards," *Authorea*, 2025. https://doi.org/10.22541/au.175192230.05564058/v1

[8] M. Rahouti, K. Xiong, and Y. Xin, "Secure software-defined networking communication systems for smart cities: Current status, challenges, and trends," *IEEE Access*, vol. 9, pp. 12083-12113, 2021. https://doi.org/10.1109/ACCESS.2020.3047996

[9] A. M. Tall, C. C. Zou, and J. Wang, "Integrating cybersecurity into a big data ecosystem," presented at the MILCOM 2021-2021 IEEE Military Communications Conference (MILCOM), IEEE, 2021.

[10] J. M. L. Velásquez, S. M. M. Monterrubio, L. E. S. Crespo, and D. G. Rosado, "SIEM-SC initial assessments: Towards a sustainable and compliant proposal for security information and event management," *International Journal of Information Security*, vol. 24, no. 5, p. 195, 2025. https://doi.org/10.1007/s10207-025-01109-w

[11] D. Fayyad, "Automated incident containment using SOAR platforms in large-scale enterprises," *Journal of Data Analysis and Critical Management*, vol. 1, no. 4, pp. 51-62, 2025. https://doi.org/10.64235/x3engj56

[12] Ismail *et al.*, "Toward robust security orchestration and automated response in security operations centers with a hyper-automation approach using agentic artificial intelligence," *Information*, vol. 16, no. 5, p. 365, 2025. https://doi.org/10.3390/info16050365

[13] N. Sun *et al.*, "Cyber threat intelligence mining for proactive cybersecurity defense: A survey and new perspectives," *IEEE Communications Surveys & Tutorials*, vol. 25, no. 3, pp. 1748-1774, 2023. https://doi.org/10.1109/COMST.2023.3273282

[14] H. V. Nguyen and T. Uehara, "Multilayer action representation based on MITRE ATT&CK for automated penetration testing," *Journal of Information Processing*, vol. 31, pp. 562-577, 2023. https://doi.org/10.2197/ipsjjip.31.562

[15] J. Henriques, F. Caldeira, T. Cruz, and P. Simões, "A survey on forensics and compliance auditing for critical infrastructure protection," *IEEE Access*, vol. 12, pp. 2409-2444, 2024. https://doi.org/10.1109/ACCESS.2023.3348552

[16] Z. A. Bilal and M. Kalamir, "Strategic approaches to cybersecurity: The role of the security operations center," *Croatian Regional Development Journal*, vol. 6, no. 2, pp. 1-12, 2025. https://doi.org/10.2478/crdj-2025-0007

[17] A. Chandwani, S. Dey, and A. Mallik, "Cybersecurity of onboard charging systems for electric vehicles—Review, challenges and countermeasures," *IEEE Access*, vol. 8, pp. 226982-226998, 2020. https://doi.org/10.1109/ACCESS.2020.3045367

[18] M. Gupta, M. Abdelsalam, S. Khorsandroo, and S. Mittal, "Security and privacy in smart farming: Challenges and opportunities," *IEEE Access*, vol. 8, pp. 34564-34584, 2020. https://doi.org/10.1109/ACCESS.2020.2975142

[19] V. J. M. López, M. S. M. Monterrubio, S. L. E. Crespo, and G. D. Rosado, "Systematic review of SIEM technology: SIEM-SC birth," *International Journal of Information Security*, vol. 22, no. 3, pp. 691-711, 2023. https://doi.org/10.1007/s10207-022-00657-9

[20] M. Lee, J. Jang-Jaccard, and J. Kwak, "Novel architecture of security orchestration, automation and response in internet of blended environment," *Computers, Materials & Continua*, vol. 73, no. 1, pp. 199–223, 2022. https://doi.org/10.32604/cmc.2022.028495

[21] B. E. Strom, A. Applebaum, D. P. Miller, K. C. Nickels, A. G. Pennington, and C. B. Thomas, *MITRE ATT&CK: Design and philosophy*. Bedford, MA: MITRE Corporation, 2018.

[22] S. De, "Security threat analysis and prevention towards attack strategies," in *Cyber Defense Mechanisms: Security, Privacy, and Challenges*, G. Kumar, D. K. Saini, and N. H. H. Cuong, Eds. Boca Raton, FL: CRC Press, 2020, pp. 1–22.

[23] V. H. U. Eze, C. N. Ugwu, and I. C. Ugwuanyi, "A study of cyber security threats, challenges in different fields and its prospective solutions: A review," *INOSR Scientific Research*, vol. 9, no. 1, pp. 13-24, 2023.

[24] J. Baldwin, O. M. K. Alhawi, S. Shaughnessy, A. Akinbi, and A. Dehghantanha, "Emerging from the cloud: A bibliometric analysis of cloud forensics studies," in *Cyber Threat Intelligence*, A. Dehghantanha, M. Conti, and T. Dargahi, Eds., Advances in Information Security, vol. 70. Cham, Switzerland: Springer, 2018, pp. 311–331. https://doi.org/10.1007/978-3-319-73951-9_16

[25] A. Mohsin, H. Janicke, A. Ibrahim, I. H. Sarker, and S. Camtepe, "A unified framework for human AI collaboration in security operations centers with trusted autonomy," *arXiv preprint arXiv:2505.23397*, 2025. https://doi.org/10.48550/arXiv.2505.23397

[26] S. Srinivas *et al.*, "AI-Augmented SOC: A survey of LLMs and agents for security automation," *Journal of Cybersecurity and Privacy*, vol. 5, no. 4, p. 95, 2025. https://doi.org/10.3390/jcp5040095

[27] A. M. Winkler and P. Sharma, "Proactive threat detection in enterprise systems using Wazuh: A MITRE ATT&CK evaluation," *Computers & Security*, vol. 159, p. 104702, 2025. https://doi.org/10.1016/j.cose.2025.104702

[28] W. Abbass, Z. Bakraouy, A. Abou El Kalam, L. El Haloui, A. Baina, and M. Bellafkih, "Building a proactive cybersecurity Defense: Detection and analysis strategies," presented at the 2025 International Conference on Intelligent Systems: Theories and Applications (SITA), IEEE, 2025.

[29] Attaullah, A. Sufyan, M. Mujeeb-Ur-Rehman, B. Noreen, and S. Amin, "Trends, capabilities, and challenges in modern cyber defense: A systematic review of detection and response technologies," *Spectrum of Engineering Sciences*, vol. 4, no. 1, pp. 464-503, 2026. https://doi.org/10.5281/zenodo.18346067

[30] S. G. Virk, J. Iqbal, A. Ali, A. R. Mahmud, I. Rashid, and T. Hanif, "Advancing security operations centers: Modern use cases, MITRE ATT&CK integration, and coverage optimization in 2025," *Journal of Computing & Biomedical Informatics*, vol. 9, no. 2, pp. 1-14, 2025.

[31] D. Fayyad, "Enhancing incident response efficiency through automated dynamic threat case generation," *International Journal of Technology, Management and Humanities*, vol. 11, no. 4, pp. 26-37, 2025. https://doi.org/10.21590/ijtmh.11.04.04

[32] J. Roy and S. K. Singh, "AgentSOC: A multi-layer agentic AI framework for security operations automation," in *2026 IEEE 5th International Conference on AI in Cybersecurity (ICAIC)*, Houston, TX, USA, 2026, pp. 1-7. https://doi.org/10.1109/ICAIC67076.2026.11395783

[33] M. Masunda, "Adaptive threat intelligence systems for real-time detection and mitigation of sophisticated cyber attacks in enterprise networks," *International Journal of Advance Research Publication and Reviews*, vol. 2, no. 5, pp. 470-491, 2025.

[34] T. K. Chowdhury, "AI-powered deep learning models for real-time cybersecurity risk assessment in enterprise IT systems," *ASRC Procedia: Global Perspectives in Science and Scholarship*, vol. 1, no. 1, pp. 675-704, 2025. https://doi.org/10.63125/137k6y79

[35] O. M. Abdelaziz and H. K. Aslan, "A unified security operations center and zero trust architecture framework for adaptive IoT threat mitigation," presented at the 2025 International Mobile, Intelligent, and Ubiquitous Computing Conference (MIUCC), IEEE, 2025.

[36] Z. Aradi, S. Bottyán, E. Kail, E. Rigó, and A. Bánáti, " Metrics-driven evaluation and optimization of honeypots: Toward standardized measures of deception effectiveness," *Acta Polytechnica Hungarica*, vol. 22, no. 12, pp. 295–313, 2025. https://doi.org/10.12700/APH.22.12.2025.12.19

[37] P. R. Rajgopal, "AI-optimized SOC playbook for ransomware investigation," *International Journal of Data Science and Machine Learning*, vol. 5, no. 2, pp. 41–55, 2025. https://doi.org/10.55640/ijdsml-05-02-04

[38] J. Kuforiji, "Digital forensics and incident response (DFIR) automation: Leveraging AI to accelerate breach investigation, evidence collection, and cyberattack mitigation," *Journal of Data Analysis and Critical Management*, vol. 1, no. 4, pp. 1-19, 2025. https://doi.org/10.64235/tsvfvz27