

Article

AI-Driven Threat Detection and Automated Incident Response for Securing Cloud Workloads

Anton Chagovec ¹, Teodora Bakardjieva ¹, Antonina Ivanova ^{1,*}, Fatima Sapundzhi ^{1,2,*}, Veselina Spasova ^{1,3} and Andriana Ivanova ¹

¹ Department of Computer Science, Varna Free University “Chernorizets Hrabar”, 84 Yanko Slavchev Street, Chaika Resort, 9007 Varna, Bulgaria; 213010037@vfu.bg (A.C.); bakardjieva@vfu.bg (T.B.); vspasova@vfu.bg (V.S.); andriana.ivanova@vfu.bg (A.I.)

² Department of Communication and Computer Engineering, South-West University “Neofit Rilski”, 66 Ivan Mihailov Street, 2700 Blagoevgrad, Bulgaria

³ Department of Informational and Communication Technology, Faculty of Engineering, Nikola Vaptsarov Naval Academy, 73 Vasil Drumev, 9002 Varna, Bulgaria

* Correspondence: antonina.ivanova@vfu.bg (A.I.); sapundzhi@swu.bg (F.S.)

Abstract

The increasing adoption of cloud computing has expanded organizational attack surfaces and created additional opportunities for identity abuse, ransomware operations, data exposure, and configuration-related security incidents. Conventional monitoring environments based primarily on static rules, fragmented telemetry, and manual triage often struggle to prioritize high-severity incidents in real time. This study evaluates the operational impact of an integrated AI-augmented cloud-native SIEM/XDR/SOAR architecture for cloud threat detection and automated incident response. A sequential mixed-methods comparative case study was conducted across two enterprise-style security environments: an AI-augmented architecture combining cloud-native SIEM, XDR telemetry unification, behavioral analytics, AI-assisted correlation, generative-AI analyst support, and SOAR automation, and a conventional baseline environment based on manual triage and signature-based controls. Three attack scenarios were analyzed: phishing-led account takeover, multi-stage ransomware, and shadow-IT data exfiltration. The AI-augmented architecture reduced mean time to triage from 17.4 h in the conventional baseline to 10.7 min and enabled ransomware containment in under five minutes through pre-configured automated response playbooks. The results also showed improved prioritization of high-severity incidents, reduced analyst review burden, and a high automated closure rate. The findings provide operational evidence for the evaluated security architecture. Limitations include single-vendor dependency, non-equivalent false-positive classification mechanisms, proprietary model internals, calibration requirements, and detection gaps involving legitimate third-party services and password-protected content.



Academic Editors: George Drosatos, Nunzio Cennamo and Stefano Toldo

Received: 18 May 2026

Revised: 25 June 2026

Accepted: 26 June 2026

Published: 29 June 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: cloud security; SIEM; XDR; SOAR; artificial intelligence; automated incident response; Security Operations Centre; behavioral analytics; incident correlation; ransomware; cloud-native security architecture

1. Introduction

Cloud computing is increasingly used as the primary infrastructure model in enterprise IT. Gartner forecasts global public cloud end-user spending to reach USD 723 billion in 2025, with Infrastructure as a Service (IaaS) and Platform as a Service (PaaS) continuing

to grow [1]. The expansion of cloud services increases the number of externally accessible systems, identities, APIs, SaaS applications, and distributed workloads requiring continuous monitoring and protection. SentinelOne reports that 80% of organizations have experienced increased attack activity in cloud environments, with data breaches, lateral movement, and cryptomining among the most frequently reported incident categories [2]. Organizations also face growing exposure to credential compromise, ransomware, data exfiltration, insecure APIs, and cloud misconfigurations.

Conventional security mechanisms based on static rules, signature matching, fragmented telemetry, and manual triage encounter limitations in large-scale cloud environments characterized by distributed infrastructure, dynamic workloads, and heterogeneous telemetry sources. Attack techniques based on credential compromise, zero-day exploitation, and multi-stage intrusion chains frequently bypass fixed detection rules. Large alert volumes increase analyst workload and contribute to alert fatigue, where relevant incidents may be delayed during triage and investigation processes [3]. Security Operations Centres (SOCs) must process high volumes of telemetry, prioritize critical incidents, reduce alert noise, and respond to complex attack chains within increasingly constrained operational timeframes.

The operational effectiveness of integrated AI-assisted security architectures in cloud environments remains an important research problem. Operational evaluation of AI-assisted security architectures encompasses anomaly detection, incident triage, prioritization, correlation, containment, and analyst workload reduction within enterprise security operations.

AI and ML techniques are increasingly incorporated into cloud security monitoring platforms to support behavioral analysis, anomaly detection, and automated incident correlation. Two technology categories are central to current enterprise cloud security operations: next-generation Security Information and Event Management (SIEM) systems and Extended Detection and Response (XDR) platforms. Existing studies and industry reports describe the operational capabilities of these platforms. Comparative empirical evaluations across multiple attack scenarios and conventional baseline environments remain limited [4].

This study examines the operational effectiveness of an integrated AI-augmented SIEM/XDR/SOAR architecture in cloud environments. The evaluation focuses on architecture-level performance indicators, including triage time, incident closure time, incident correlation, automated closure rates, analyst review workload, and response behavior across different attack scenarios.

Research Gap and Objectives

Empirical evaluations of AI-augmented cloud security monitoring systems remain limited in several respects. Many studies focus on algorithm-level detection accuracy, isolated attack scenarios, or vendor-specific platform descriptions. Operational effects of integrated AI-assisted correlation, XDR telemetry unification, SOAR automation, and cloud-native SIEM engineering on SOC workflows under production-style conditions are examined less frequently. This study evaluates operational metrics including triage time, closure time, incident correlation ratio, observed false-positive activity, true-positive classification, automated closure rate, and ransomware containment time.

The scope of the study is defined by the following research objectives:

- RO1: To examine the effect of an AI-augmented SIEM/XDR/SOAR architecture on detection, triage, and response time across multiple cloud attack scenarios.
- RO2: To evaluate observed false-positive activity and analyst review burden associated with behavioral analytics and AI-assisted event correlation in comparison with a conventional rule-based environment.

- RO3: To analyze the effect of AI-assisted incident correlation, XDR telemetry unification, and automated response mechanisms on analyst workload and incident handling efficiency.
- RO4: To identify operational limitations and governance considerations associated with AI-supported and SOAR-based automated incident response in enterprise cloud environments.

The analysis focuses on SIEM/XDR/SOAR workflows at the architectural level:

- A comparative architecture-level evaluation of an AI-augmented SIEM/XDR/SOAR environment and a conventional security operations baseline across three cloud attack scenarios.
- A production-oriented measurement model using triage time, closure time, incident correlation ratio, observed false-positive activity, true-positive classification, automated closure rate, and ransomware containment time.
- An analysis of AI-assisted correlation, XDR telemetry unification, generative-AI analyst support, and SOAR playbooks in phishing-led account takeover, ransomware, and shadow-IT exfiltration scenarios.
- Identification of operational limitations related to behavioral model calibration, legitimate third-party services, encrypted and password-protected content, vendor dependence, proprietary model internals, and governance of automated response.

The analysis focuses on the operational behavior of an integrated security architecture using production-style incident data, scenario-based analysis, and SOC-oriented operational metrics. The reported findings represent descriptive operational evidence for the combined architecture and do not establish causal effects attributable to an individual AI module.

2. Background and Related Work

2.1. Cloud Security Threat Landscape

Cloud-related risks such as misconfiguration, weak identity and access management, insecure APIs, and insufficient monitoring are widely documented in cloud security literature and industry threat reports [5–8]. Their significance is amplified by the distributed and dynamic nature of cloud infrastructures, where configuration changes, service interactions, and identity management processes occur at large scale.

Compromised credentials remain a common factor in cloud incidents, and the ENISA Threat Landscape report identifies ransomware and data theft among the most persistent threats affecting cloud-connected infrastructure [6,7]. Akamai reports that 84% of surveyed security professionals experienced API-related security incidents during 2024 [8]. The global cybersecurity workforce shortage continues to affect operational response capacity. ISC² estimates a shortage of 4.76 million cybersecurity professionals worldwide [9].

2.2. SIEM Evolution and AI Integration

SIEM platforms were initially developed for log aggregation, event correlation, and compliance monitoring. Earlier systems relied primarily on predefined correlation rules and signature-based detection mechanisms. Cloud-native SIEM architectures increasingly incorporate behavioral analytics, anomaly detection, and ML-based correlation mechanisms to process large volumes of telemetry data from distributed infrastructure. This evolution has expanded the operational role of SIEM platforms beyond log aggregation and compliance monitoring. Gartner defines modern SIEM platforms as systems that aggregate and analyze event data from cloud and on-premises infrastructure [10].

User and Entity Behavior Analytics (UEBA) extends SIEM functionality through behavioral profiling and anomaly detection based on historical activity patterns. AI-

assisted correlation engines aggregate related events from multiple telemetry sources into contextual incidents requiring analyst attention. CrowdStrike describes AI-powered SIEM environments as capable of reducing alert noise and improving prioritization of high-severity events [11].

2.3. XDR Platforms and AI Integration

Extended Detection and Response (XDR) platforms extend the scope of endpoint-focused monitoring by integrating telemetry from endpoints, networks, cloud workloads, identity systems, email infrastructure, and SaaS applications. Telemetry normalization and orchestration mechanisms for scalable security monitoring across distributed multi-cloud infrastructures have been examined in recent studies [12,13]. The integration of multiple telemetry sources enables investigation of security events across different infrastructure layers and supports a unified operational view of attack activity. Palo Alto Networks defines XDR as a technology combining multiple telemetry sources for threat detection and response [14]. Gartner describes XDR platforms as integrated environments supporting automated investigation and cross-domain event correlation [15].

Behavioral analysis and ML-assisted anomaly detection represent central operational components in current XDR architectures. Reinforcement learning approaches are being investigated for adaptive policy optimization and automated response adjustment [16]. Large language models (LLMs) are increasingly incorporated into investigation workflows for contextual summarization, natural-language querying, and response orchestration support.

2.4. Automated Incident Response

Security Orchestration, Automation, and Response (SOAR) capabilities are increasingly integrated into SIEM and XDR platforms to support automated execution of response procedures following threat detection events [17]. Cloud security automation architectures combining AI-assisted detection, orchestration workflows, and automated response coordination are discussed in recent enterprise cloud security studies [18]. AI-assisted SOC environments have been associated with improved incident prioritization and reduced operational workload during large-scale security monitoring [19]. Recent empirical evaluations of ATT&CK-aligned SIEM environments report improvements in detection coverage, alert prioritization, and operational SOC workflows in enterprise monitoring scenarios [20]. Comparative empirical evaluations across multiple attack scenarios, heterogeneous infrastructures, and conventional baseline environments remain scarce [21].

2.5. Theoretical Framework

This study is based on the Detect–Respond–Recover model of the NIST Cybersecurity Framework [22] and the Zero Trust Architecture principles defined in NIST SP 800-207 [23], which require continuous verification of entities independent of network location. Recent cloud security research also examines hybrid deep-learning approaches operating within Zero-Trust architectures for real-time threat detection in distributed cloud environments [24]. Dynamic capabilities theory [25] provides the conceptual basis for examining AI integration in SIEM/XDR environments through three operational dimensions: real-time behavioral anomaly detection across cloud telemetry, automated response orchestration through security playbooks, and adaptive ML model adjustment based on incident feedback. The evaluation metrics—MTTD, MTTR, false-positive rate, and incident-correlation precision—correspond to SOC performance indicators commonly used in MITRE ATT&CK-oriented detection engineering practice [4].

3. Materials and Methods

3.1. Research Design

This study employs a sequential mixed-methods design [26]. This design was selected because the research problem required both quantitative operational measurement and qualitative interpretation of scenario-specific behavior. The quantitative phase measures detection and response performance across multiple attack scenarios in two operational environments. The qualitative phase examines the contextual behavior of the evaluated architecture during representative security incidents. The qualitative phase interprets selected production incident cases to explain how AI-assisted correlation, XDR telemetry unification, generative-AI analyst support, and SOAR-based automation affected handling in practice.

The evaluation was conducted in production-style security environments, which precluded controlled laboratory experiments, formal randomization, and identical replay of attack traces across environments. The reported metrics represent descriptive operational indicators. The research design combines operational performance measurement with incident-level analysis. Formal causal attribution of individual architectural components remains outside the scope of the study.

The research design combines three layers of evidence: (i) quantitative incident metrics extracted from SIEM/XDR and ticketing records; (ii) scenario-based analysis of three representative cloud attack chains; and (iii) qualitative interpretation of operational limitations and governance requirements.

3.2. Evaluation Environments

Environment A represents the AI-augmented architecture under primary evaluation. It comprises a cloud-native, data-lake-based SIEM with integrated UEBA and ML-assisted correlation; an XDR layer providing unified telemetry ingestion from endpoints, cloud workloads, identity systems, email, network, and SaaS sources; SOAR-integrated playbooks for pre-approved containment actions; and a generative-AI investigation assistant supporting natural-language queries, incident summarization, timeline reconstruction, and response coordination. The evaluation focused on the operational behavior of this integrated architecture within a deployed production-style security environment.

The modeling and correlation layer is described at a functional level because the exact internal algorithms of the vendor-managed Fusion engine are proprietary and not fully accessible to the authors. Observable functions included behavioral baselining, identity risk scoring, anomaly detection, entity enrichment, multi-source event correlation, and contextual incident fusion. The evaluation is performed at the functional level and does not extend to algorithm-level replicability or separate assessment of clustering, supervised classification, or specific anomaly detection algorithms.

The generative-AI assistant supported investigation and response activities. Decision-making functions remained part of the analyst workflow. Analysts interacted with it through natural-language queries regarding incident context, affected entities, correlated alerts, user activity, and attack timeline reconstruction. The assistant generated summaries and contextual explanations, but high-impact response actions such as account suspension, device isolation, and application blocking were executed by pre-configured SOAR playbooks according to predefined severity thresholds, correlation rules, and approved operational procedures.

Figure 1 illustrates the generalized automated incident response architecture deployed in Environment A, from multi-source telemetry ingestion and XDR unification through AI-assisted correlation, generative-AI analyst support, SOAR-based containment actions, and analyst escalation. The generative-AI assistant supports investigation and coordination,

whereas high-impact response actions remain governed by pre-configured SOAR logic and analyst validation procedures.




Figure 1. AI-augmented SIEM/XDR/SOAR architecture for incident detection and response.

The comparison between Environment A and Environment B examines operational characteristics at the architecture level. The two environments differ in the presence of AI-assisted analytics, telemetry integration, automation maturity, identity protection, endpoint coverage, and response orchestration. The reported improvements reflect the operational effect of the integrated AI-enhanced SIEM/XDR architecture and do not establish the causal contribution of a single AI module.

Environment B: Conventional Baseline

Environment B represents a conventional security monitoring baseline: legacy signature-based endpoint protection, on-premises Active Directory without cloud identity integration, perimeter firewall rules, limited cloud threat protection capabilities, and manual triage and response workflows. This configuration is representative of organizations that have not yet adopted AI-augmented SIEM/XDR and SOAR capabilities. Product-specific identifiers and selected operational details were generalized to preserve organizational confidentiality and reduce vendor-specific bias.

3.3. Attack Scenarios

The scenarios were extracted from records of production incidents and operational security exercises. Sensitive identifiers, tenant information, user data, infrastructure details, and product-specific operational information were anonymized or aggregated. Three attack chains were selected to represent prevalent and impactful threat categories targeting cloud workloads, consistent with ENISA threat landscape categories and cloud security literature [6,7].

Scenario 1—Phishing-Led Account Takeover (S1): A phishing email was delivered to a target user via a legitimate third-party service, leading to credential theft. The attacker subsequently accessed cloud services from an anomalous geographic location, performed lateral movement within the cloud tenant, and exfiltrated data from a SaaS application.

Scenario 2—Multi-Stage Ransomware (S2): Initial access was achieved via a compromised personal device belonging to a privileged administrator over VPN. The attack chain included credential theft tool deployment (Mimikatz-equivalent), defense evasion, command-and-control establishment (Cobalt Strike-equivalent), and ransomware execution targeting cloud-connected server workloads, spanning 220+ entities across two servers.

Scenario 3—Shadow-IT Data Exfiltration (S3): A user transferred sensitive organizational data to an unsanctioned cloud application using a password-protected document distributed through a legitimate document-signing service. The use of encrypted content and trusted third-party infrastructure bypassed standard DLP inspection and delayed identification during the initial transfer stage (Tables 1 and 2).

Table 1. Attack scenarios, telemetry sources, and evaluated AI/XDR/SOAR functions.

Scenario	Definition	Initial Vector	Main Stages of the Attack	Telemetry Sources	AI/XDR/SOAR Feature Assessment
S1:	Account takeover via phishing	Phishing via a legitimate third-party service	Credential theft; anomalous login; lateral movement; SaaS data access	Email, Identity, SaaS, Cloud Activity	Identity Risk Assessment; Multi-Source Correlation; Session Shutdown and Password Reset
S2:	Multi-stage ransomware	Compromised privileged device via VPN	Credential theft; defense evasion; C2; ransomware execution	Endpoint, VPN, Identity, Cloud-Connected Servers	Multi-Stage Fusion Correlation; Device Isolation; Account Deactivation; Enrichment
S3:	Shadow-IT data exfiltration	Unsanctioned cloud application and password-protected document	Sensitive data transfer; impossible travel; unusual download patterns	CASB/SaaS, Identity, Network, Cloud Application Logs	Behavioral Anomaly Detection; App Blocking; Analyst Escalation

Table 2. Anonymized MITRE ATT&CK mapping and kill-chain timeline for the evaluated scenarios.

Scenario	Definition	Initial Vector	Main Stages of the Attack	Telemetry Sources	AI/XDR/SOAR Feature Assessment
S1:	Initial access	Phishing message delivered through a legitimate third-party service with protected content	T1566.001/T1566.002	Email, third-party service logs	Email filtering partially bypassed; later identity-based detection
S1	Credential access/valid account use	Credential theft followed by anomalous cloud sign-in from atypical geography	T1078; T1110.003	Identity, sign-in, SaaS activity	Identity risk assessment; unfamiliar sign-in and impossible-travel correlation
S1	Lateral movement/collection	Access to SaaS resources and cloud tenant activity outside normal behavior	T1087; T1530	SaaS, cloud activity logs	Multi-source correlation; session shutdown; password reset
S2:	Initial access	Compromised privileged personal device connecting through VPN	T1078; T1133	VPN, identity, endpoint telemetry	XDR correlation; privileged account risk escalation
S2	Credential access	Credential dumping tool behavior on cloud-connected server workloads	T1003.001	Endpoint, server, identity telemetry	Fusion correlation; enrichment; high-severity incident creation

Table 2. *Cont.*

Scenario	Definition	Initial Vector	Main Stages of the Attack	Telemetry Sources	AI/XDR/SOAR Feature Assessment
S2	Command and control/execution	C2-like activity followed by ransomware execution across connected workloads	T1071.001; T1486	Endpoint, network, cloud-connected servers	Automated device isolation; account deactivation; war-room escalation
S3:	Exfiltration preparation	Password-protected document and unsanctioned cloud service use	T1567.002; T1020	CASB/SaaS, network, cloud application logs	Initial content inspection limited; behavior-based detection later
S3	Exfiltration/anomalous behavior	Impossible travel and unusual download volume patterns associated with shadow IT use	T1537; T1567.002	Identity, CASB, SaaS logs	Behavioral anomaly detection; app blocking; analyst escalation

The selected scenarios represent complementary operational contexts for evaluating AI-assisted detection, incident correlation, and automated response capabilities. The following section defines the metrics and analytical approach used to compare the two security environments.

3.4. Metrics and Analytical Approach

The following metrics were collected for both environments across all three scenarios. Based on the operational records that did not consistently separate detection, confirmation, and response timestamps, this study used mean time to triage (MTT) and mean time to closure (MTC) as operational metrics for SOC. They are treated as practical workflow metrics and not as direct substitutes for the standard metrics of MTTD and MTTR (Table 3):

- Mean Time to Triage (MTT): Elapsed time from incident registration to first analyst or automated review action.
- Mean Time to Closure (MTC): Elapsed time from initial triage to full incident resolution.
- Incident Correlation Ratio (ICR): Ratio of raw alerts to AI-correlated contextual incidents.
- Observed False-Positive Activity (FPA): Proportion of reviewed incidents classified as FalsePositive upon investigation. Because the two environments used different classification mechanisms, this value serves as an operational indicator and should not be treated as a directly equivalent statistical metric.
- True-Positive Rate (TPR): Proportion of reviewed incidents confirmed as genuine threats.
- Automated Closure Rate (ACR): Proportion of incidents closed automatically without analyst intervention.

Data were extracted from SIEM incident management tables (SecurityIncident) during a 30-day production measurement period. Fusion/Multi-stage AI-correlated incidents were identified using title pattern matching. Incident classifications (TruePositive, FalsePositive, BenignPositive, Undetermined) were recorded upon closure and verified through routine SOC incident review procedures. Environment B metrics were obtained from manual incident logs and ticketing records covering an equivalent period. Environment B did not provide automated classification metadata. False-positive activity is reported as an approximate operational indicator and should not be interpreted as a statistically equivalent measure.

The comparison evaluates the operational impact of an integrated, AI-enhanced SIEM/XDR security architecture against a conventional monitoring baseline. The results

characterize operational performance at the architecture level and reflect the combined effect of the architectural components operating within the evaluated environment.

The two evaluated environments differed in architecture, telemetry coverage, automation maturity, tool integration, and operational constraints, and no formal inferential testing was applied. The comparison was based on descriptive operational indicators, percentage changes, and effect ratios. Operational performance was examined at the architectural level under production-style conditions. The reported metrics characterize the behavior of the evaluated environments within their operational context. Evaluation of the causal contribution of individual components would require more controlled study designs.

To improve comparability, only incidents belonging to equivalent operational categories were included in the comparison: identity, endpoint, ransomware-related activity, cloud application abuse, and suspicious data transfer. Informational events and automatically generated low-risk notifications without analyst relevance were excluded from the Environment B false-positive estimate. For Environment A, classification labels were obtained from SIEM/XDR incident records. For Environment B, classifications were reconstructed from analyst ticket notes and closure categories. Because the two environments used different classification mechanisms, false-positive activity is reported as an operational indicator.

Table 3. Evaluation metrics and calculation method.

Metric	Definition	Calculation
MTT	Time from incident registration to first analyst/automated action	First triage timestamp-incident timestamp
MTC	Time from triage to closure	Closure timestamp-triage timestamp
ICR	Alert aggregation ratio	Total raw alerts/AI-correlated incidents
FPA	Observed false-positive activity indicator	FalsePositive incidents/reviewed incidents (classification-source dependent)
TPR	Confirmed threat share	True positives/reviewed incidents
ACR	Automated closure rate	Auto-closed incidents/total incidents

3.5. Environment Configuration Summary

The evaluated environments differed substantially in architecture, security tooling, automation capabilities, licensing scope, and operational integration, representing two distinct maturity levels of enterprise cloud security monitoring that reflect deployment configurations commonly observed in contemporary organizational practice. Table 4 summarizes the main configuration parameters used in the comparative evaluation.

Product names, licensing labels, and operational identifiers were generalized to preserve confidentiality and to reduce vendor bias.

The product names and licensing labels were summarized to preserve confidentiality and reduce bias toward a specific vendor.

Table 4. Environment Configurations and Evaluation Parameters.

Parameter	Environment A (AI-Augmented)	Environment B (Conventional Baseline)
SIEM Platform	Cloud-native data-lake SIEM with ML/UEBA and AI-assisted multi-stage correlation engine (like Fusion)	None integrated SIEM; manual log review only
XDR Layer	Integrated XDR layer covering endpoint, cloud, identity, e-mail and network telemetry (Vendor-agnostic XDR)	Siloed point tools: perimeter firewall and legacy endpoint protection (Cisco ASA (Cisco Systems, San Jose, CA, USA)/Symantec SEP (Broadcom Inc., Palo Alto, CA, USA))

Table 4. Cont.

Parameter	Environment A (AI-Augmented)	Environment B (Conventional Baseline)
Endpoint Protection	Cloud-delivered endpoint protection with behavioral and static AI engines (Next-gen AI/ML)	Legacy signature-based endpoint protection (like Symantec SEP)
Identity Protection	Cloud identity protection with ML risk scoring and conditional access policies	On-premises identity management only; no cloud risk scoring
Email Security	AI-assisted e-mail threat protection with URL rewriting, attachment sandboxing, and post-delivery remediation.	Basic e-mail filtering only
Gen-AI Assistant	Generative AI investigation assistant supporting natural-language query, summarization, and playbook coordination (like GPT)	Not available
Response Automation	SOAR-integrated playbooks for auto-isolation, account locking, and domain/application blocking	Manual only; no automation
Licensing	Enterprise-grade integrated security licensing (like E5-equivalent)	Limited cloud productivity/security licensing
Measurement Period	30 days (production)	30 days (production)

4. Results

4.1. Measurement Model Validation: Alert Correlation Performance

To evaluate alert correlation and prioritization performance in Environment A, incident classification outcomes were analyzed across the full 30-day period (Figure 2). The Fusion/Multi-stage correlation engine generated 66 AI-correlated incidents from a base of 2694 total logged alerts, corresponding to an incident correlation ratio (ICR) of approximately 41:1, indicating that each AI-correlated case aggregated an average of 41 raw alerts into a contextual incident.

The majority of Fusion-correlated incidents were classified as BenignPositive or TruePositive. Of the 66 Fusion incidents, 54 (81.8%) were classified as BenignPositive or TruePositive and required analyst attention. Five incidents (7.6%) were confirmed as TruePositive active threats. Forty-nine incidents (74.2%) were classified as BenignPositive suspicious but expected activity. The remaining 12 incidents (18.2%) were closed without definitive classification. Fusion correlation was used to aggregate, prioritize, and contextualize related security events for analyst review.

In this study, incidents with benign positive characteristics (BenignPositive) are not treated as false alarms. They represent suspicious but expected activities that require validation by the analyst. The Fusion Engine is primarily evaluated as a prioritization and correlation mechanism. Evaluation focuses on incident aggregation, prioritization, contextualization, and threat classification performance.

BenignPositive incidents represented the largest share of Fusion-correlated incidents. The correlated incident population was dominated by suspicious but expected activity requiring analyst validation.

For the full incident population, the XDR layer automatically closed 89.42% of incidents ($n = 2341$ of 2628 non-Fusion incidents) without analyst intervention, corresponding to an automated closure rate (ACR) of 89.4%. The analyst-reviewed population comprised 277 incidents, of which 258 (93.1%) were confirmed BenignPositive and 19 (6.9%) were

confirmed TruePositive—a true-positive rate (TPR) of 0.73% of all incidents, representing the confirmed genuine threat caseload for human analysts over 30 days.




Figure 2. Classification outcomes of AI-correlated Fusion incidents in Environment A during the 30-day measurement period ($n = 66$).

4.2. Structural Model Results: Severity Distribution

The incident classification results and severity distribution observed during the 30-day evaluation period in Environment A are summarized in Table 5, providing a disaggregated view of AI-correlated and non-AI incidents across all severity levels and outcome categories.

Table 5. Incident Classification and Severity Distribution (30-Day Period, Environment A, $N = 2694$).

Metric	AI-Fusion Incidents ($n = 66$)	Non-AI Incidents ($n = 2628$)	Total ($N = 2694$)
Auto-closed by XDR (Undetermined)	12 (18.2%)	2341 (89.4%)	2353
BenignPositive (analyst confirmed)	49 (74.2%)	258 (9.85%)	307
TruePositive (confirmed threat)	5 (7.6%)	19 (0.73%)	24
High severity	49 (74.2%)	200 (7.6%)	249
Medium severity	12 (18.2%)	266 (10.2%)	278
Low severity	1 (1.5%)	48 (1.8%)	49
Informational	4 (6.1%)	2104 (80.4%)	2108
Incident Correlation Ratio (ICR)	~41:1 (alerts to incidents)	1:1 (no aggregation)	-

Fusion incidents demonstrated a 74.2% High-severity rate versus 7.6% for non-AI incidents, indicating prioritization of higher-severity incidents by the AI correlation engine (ratio: 9.7 $\times$).

The severity distribution indicates that the Fusion correlation layer preferentially surfaced higher-severity incidents from the broader alert population: High-severity cases increased from 7.6% in the non-AI incident population to 74.2% among AI-correlated cases, representing a 9.7-fold enrichment. This behavior suggests that the correlation layer supported prioritization of incidents requiring analyst attention and is consistent with RO3.

4.3. Observed False-Positive Activity and Analyst Review Burden (RO2)

Lower observed false-positive activity and a reduced analyst review burden were recorded in the environment that included behavior-profiling mechanisms in the XDR layer. The reported values originate from different classification mechanisms and are not interpreted as strictly equivalent statistical measurements.

In Environment A, the proportion of analyst-reviewed incidents classified as False-Positive was approximately 10.6% during the measurement period, based on automated SIEM/XDR incident metadata. In Environment B, the false-positive activity estimate of approximately 62% was reconstructed from manual ticketing records and analyst closure notes. Because Environment B did not provide automated classification metadata, this value is reported as an estimated operational indicator.

BenignPositive incidents represented suspicious but expected activity requiring analyst validation. False alarms were reserved for incidents without security relevance after review. They represented suspicious but expected activity requiring analyst validation. FalsePositive was reserved for incidents determined to have no security relevance after review. The comparison describes differences in alert prioritization and analyst review burden. The reported values originate from two different classification mechanisms and constitute a descriptive assessment of the evaluated environments.

4.4. Detection and Response Timing (RO1, RO4)

Table 6 presents the comparative detection and response metrics for the evaluated environments, including mean triage time, incident closure time, observed false-positive activity, and ransomware containment performance.

Table 6. Detection and Response Time Comparison: Environment A vs. Environment B.

Metric	Environment A (AI-Augmented)	Environment B (Conventional)
Mean Time to Triage (MTT)	10.7 min	17.4 h
Mean Time to Closure (MTC)	23.8 min	35.9 min *
MTT Improvement Factor	97.6× faster	- (baseline)
Ransomware containment time	<5 min (automated playbook)	~90+ minutes (manual)
Ransomware response improvement	18× faster	- (baseline)
Automated closure rate (ACR)	89.4% of all incidents	0% (manual only)
Observed false-positive activity (FPA)	~10.6% (automated metadata)	~62% (manual estimate)
Analyst triage time reduction	approximately 99%	- (baseline)
SLA compliance (MTT < 20 min)	Yes (10.7 min)	No (17.4 h)

* MTC in Environment B is understated as it excludes incidents not yet resolved within the measurement window. MTT and ransomware containment time provide more reliable indicators than MTC. The MTT comparison (17.4 h versus 10.7 min) represents an approximately 97.6-fold reduction in triage time.

Environment A achieved an MTT of 10.7 min, remaining below the contractual SLA threshold of 20 min. Environment B's MTT of 17.4 h substantially exceeded the same threshold, reflecting the structural bottleneck of manual triage without automated pre-classification. The measured MTT values differed by approximately 97.6-fold, from 17.4 h in Environment B to 10.7 min in Environment A.

To provide a consolidated view of the evaluation outcomes, the main findings are mapped to the corresponding research objectives. Table 7 summarizes the evidence, metrics, and interpretation associated with each objective.

Quantitative operational metrics and scenario-based analysis provided evidence for all four research objectives. The largest differences between the two environments were observed in triage efficiency, incident prioritization, alert correlation, and automated ransomware containment. Qualitative findings complement the quantitative results by describing operational limitations, governance requirements, and analyst involvement in AI-supported security operations.

Table 7. Mapping of research objectives to results and interpretation.

Research Objective	Corresponding Evidence	Primary Metrics/ Scenarios	Interpretation
RO1:	Detection and response timing improved in the AI-augmented architecture.	MTT: 10.7 min vs. 17.4 h; ransomware containment under 5 min in S2.	Integrated SIEM/XDR/SOAR workflows reduced triage and containment intervals.
RO2:	Observed false-positive activity and analyst review burden were lower in Environment A.	Environment A automated metadata; Environment B manual ticket estimate.	Findings indicate improved prioritization but are interpreted descriptively due to non-equivalent classification methods.
RO3:	AI-assisted correlation reduced alert fragmentation and enriched high-severity incidents.	ICR approximately 41:1; High-severity enrichment 9.7× among Fusion incidents.	Correlation and telemetry unification improved analyst focus on investigation-relevant cases.
RO4	Limitations and governance requirements were identified across all scenarios.	S1/S3 encrypted content and third-party service abuse; S2 high-impact automated containment.	AI-supported automation requires layered controls, thresholds, human oversight, and auditability.

4.5. Qualitative Findings: Attack Scenario Outcomes (RO4)

4.5.1. Scenario S1: Phishing-Led Account Takeover

The Fusion correlation engine correlated four distinct incident signals—malicious IP access, atypical geographic travel, unfamiliar sign-in properties, and password-spray activity—into a single High-severity contextual incident within minutes of initial attack activity. Identity protection services detected the anomalous login pattern, and automated playbooks suspended the compromised session and triggered a password reset. The phishing vector was not intercepted at email delivery because it used a password-protected document distributed through a legitimate third-party service. The document bypassed current AI-assisted email and content filtering mechanisms. Post-compromise containment was achieved rapidly, and protection was extended to other potentially targeted users through automated classifier updates.

4.5.2. Scenario S2: Multi-Stage Ransomware

The ransomware attack chain included privileged VPN access, credential theft, defense evasion, command-and-control communication, and ransomware execution against cloud-connected server workloads spanning more than 220 entities across two servers. Endpoint, identity, VPN, and cloud-server telemetry were correlated into a single High-severity incident covering the complete attack chain. Automated SOAR playbooks isolated compromised endpoints, deactivated the privileged account, and initiated incident escalation within minutes of detection. The coordinated response reduced the ransomware propagation window to less than five minutes, compared with an estimated 90+ minutes under manual procedures in Environment B, and limited the number of systems requiring manual remediation.

4.5.3. Scenario S3: Shadow-IT Data Exfiltration

Cloud application security monitoring detected impossible-travel anomalies and unusual download volume patterns consistent with data exfiltration. Automated policies blocked access to the unsanctioned application and generated an analyst alert. Behavioral monitoring identified suspicious exfiltration patterns after the initial transfer. The password-protected document distributed through a legitimate third-party service was not intercepted at ingestion, highlighting a remaining limitation in content inspection and automated classification. Layered controls and human oversight remained necessary in

social-engineering and encrypted-content scenarios. Analyst judgment continued to play a role in the investigation and response process.

5. Discussion

5.1. Interpretation of Findings

The evaluation results indicate that integrated AI-augmented SIEM/XDR/SOAR architectures can improve detection prioritization and response speed in cloud security operations. Across the evaluated attack scenarios, the AI-augmented environment demonstrated shorter triage intervals, stronger contextual alert aggregation, and faster containment of high-severity incidents. The observed operational differences were most visible during multi-stage attack activity involving correlated telemetry from cloud workloads, identity systems, endpoints, email, SaaS services, and network sources.

For RO1, the most visible operational effect was the reduction in triage time. RO2 was associated with lower observed false-positive activity and a reduced analyst review burden, taking into account the differences between automated SIEM/XDR metadata and manually reconstructed baseline ticket records. RO3 was reflected in reduced alert fragmentation through the aggregation of related signals into contextual incidents and the escalation of a greater proportion of high-severity cases for analyst review. RO4 highlighted limitations associated with the abuse of legitimate services, verification of encrypted or password-protected content, and governance of automated response actions.

Environment A aggregated related alerts from multiple telemetry sources into contextual incidents with broader investigative context [17]. Incident activity associated with endpoints, cloud services, user identities, and network events was processed within unified investigation chains. The severity distribution of Fusion-correlated incidents also showed a substantially higher proportion of High-severity cases compared with the non-AI incident population.

In the ransomware scenario, automated response orchestration reduced the interval between initial detection and containment activity. Automated isolation procedures and identity-based response actions reduced the propagation window and limited the number of systems requiring manual remediation.

Differences were also observed in analyst workload associated with false-positive investigation and repetitive triage activity. Environment B relied primarily on static rules and manual review procedures. Investigation overhead increased during larger alert volumes, and prioritization required repeated analyst intervention. Environment A incorporated behavioral baselining and contextual event correlation before escalation to analysts, reducing the volume of alerts requiring direct review.

The qualitative findings also emphasize the role of human oversight in AI-assisted security operations. Password-protected documents distributed through legitimate third-party services were not intercepted during the initial delivery stage in S1 and S3. These cases demonstrate that AI-assisted filtering, DLP inspection, and behavioral analytics may have limited visibility into encrypted or password-protected content until later behavioral anomalies emerge.

This limitation reinforces the need for layered controls, including identity protection, behavioral monitoring, cloud application security, user awareness, attachment sandboxing where technically possible, restrictions on unsanctioned applications, and escalation procedures for suspicious third-party service activity. AI-supported systems can accelerate detection and response after suspicious behavior becomes visible, but they cannot fully replace analyst judgment in ambiguous social-engineering and encrypted-content scenarios.

Behavioral analytics, SIEM/XDR integration, and SOAR automation improve SOC alarm prioritization and reduce alarm noise in modern environments. Unlike model-based intrusion detection studies that primarily look at classification accuracy, this study emphasizes operational SOC metrics and response behavior at the architectural level.

5.2. Comparison with Related Empirical Studies

The observed reduction in triage time and alert fragmentation is consistent with recent studies emphasizing the role of AI-assisted SIEM frameworks in reducing alert fatigue and improving incident response efficiency [10]. The results also align with research on SIEM trends in critical infrastructures, which identifies event correlation, prioritization, and integration of heterogeneous telemetry as central requirements for modern security monitoring [11]. Compared with MITRE ATT&CK-aligned evaluations of enterprise monitoring environments [20], the present study contributes an operational architecture-level perspective by combining incident correlation, XDR telemetry unification, and SOAR-based containment across multiple cloud attack scenarios.

Environment B represents a conventional baseline characterized by fragmented telemetry and limited automation. More advanced SOC environments may operate with broader telemetry integration, higher automation maturity, and different investigation workflows. The comparison is intended as an operational baseline assessment and is supplemented by benchmarking against related empirical studies.

5.3. Practical Implications

The findings show that integrated SIEM/XDR/SOAR environments with behavioral analytics, telemetry unification, and automated response capabilities can improve operational efficiency in enterprise cloud security monitoring. The evaluated architecture reduced the volume of incidents requiring manual review and shortened the response cycle for high-severity events. These improvements should be understood as the result of the integrated architecture and its constituent components.

Generative AI assistants were used during incident investigations to generate natural-language summaries of correlated alerts and provide contextual information related to affected entities and attack activity. In the evaluated environment, these functions reduced part of the manual effort associated with triage and multi-stage incident review. Analyst involvement remained necessary for incident validation and escalation decisions, particularly in scenarios involving ambiguous activity patterns or incomplete contextual data.

The comparison between the evaluated environments also illustrates the limitations of conventional rule-based monitoring approaches in large-scale cloud infrastructures. Manual triage and fragmented telemetry sources increased investigation time and reduced prioritization efficiency in the baseline environment.

5.4. Limitations

The evaluation was conducted within a single enterprise-style cloud environment and primarily within one vendor ecosystem. Infrastructure configurations, telemetry coverage, organizational security practices, licensing levels, XDR platforms, and LLM backends may affect the applicability of the reported results to other enterprise environments. Cross-vendor generalization was not evaluated.

The comparison does not isolate the contribution of individual AI modules, SOAR automation, XDR telemetry unification, or cloud-native engineering. An ablation analysis was not feasible in the production-style environment because disabling or isolating components would have disrupted operational security processes. Additional studies across multiple vendor stacks and comparable attack traces would provide a broader basis for comparison, including different XDR platforms and generative-AI backends.

Behavioral profiling models required an initial calibration period before stable classification performance was achieved. During this phase, anomaly detection precision was lower than during later operations. The internal algorithms used by the Fusion engine and related behavioral models are proprietary and were not independently inspectable. The modeling layer is described at a functional level, reflecting the available observability of the evaluated platform.

The measurement period was limited to 30 days and focused on a selected set of attack scenarios. Additional long-term evaluations across heterogeneous multi-cloud environments may produce different correlation and response characteristics. The use of generative-AI assistants also introduces the possibility of incorrect contextual interpretation or excessive reliance on automated recommendations during incident analysis.

The absence of formal statistical testing represents a further methodological limitation. As the evaluation was conducted within live production environments, controlled randomization and hypothesis testing were not feasible. Findings should be interpreted as exploratory and case-specific within the evaluated environment.

It is also possible that the same attack traces could not be reproduced under identical controlled conditions in both environments. This comparison reflects operational differences between two security architectures. Differences in telemetry coverage, licensing level, analyst workflow maturity, automation availability, and classification metadata may have contributed to the observed performance differences. In particular, the false-positive activity comparison is descriptive because Environment A used automated SIEM/XDR metadata, whereas Environment B relied on reconstructed manual ticketing records.

5.5. Ethical and Governance Considerations

Automated response mechanisms in production cloud environments require clearly defined governance procedures, escalation policies, and confidence thresholds. Actions such as account suspension, device isolation, or automated blocking may affect operational continuity if triggered incorrectly. The evaluated ransomware scenario demonstrated the value of rapid automated containment, but it also confirmed the need to combine automated response with analyst supervision and escalation procedures for high-impact incidents.

A risk-tiered governance model supports automated response by aligning response actions with confidence levels and operational impact. Medium-impact actions, such as temporary session termination, step-up authentication, or additional verification, require higher confidence and correlation across multiple telemetry sources. High-impact actions, such as account suspension, device isolation, or blocking of business-critical applications, should require very high-confidence multi-source correlation, predefined emergency playbooks, or human approval depending on the operational context.

Organizations should define automation thresholds through documented playbook policies, audit logging, role-based approval, rollback procedures, exception handling, periodic validation of automation rules, and post-incident review. Human-in-the-loop governance remains essential for ambiguous, high-impact, or business-critical incidents, especially when AI-generated summaries or recommendations are used during investigation.

The increasing use of AI techniques in offensive operations introduces additional challenges for cloud security monitoring. AI-assisted social engineering, adversarial ML techniques, and automated evasion mechanisms may reduce the effectiveness of existing detection models. These developments require continuous model adaptation, monitoring procedures, and periodic validation of automated response mechanisms.

6. Conclusions

This study examined the operational impact of an integrated AI-augmented SI-EM/XDR/SOAR architecture for cloud threat detection and automated incident response in comparison with a conventional security environment. The evaluation should be interpreted as a production-oriented comparative case study of two security architectures and not as a controlled laboratory benchmark of individual AI models.

In the evaluated environment, the AI-augmented architecture reduced mean time to triage, improved prioritization of high-severity incidents, strengthened alert correlation, lowered analyst review burden, and enabled rapid ransomware containment through pre-configured automated response playbooks. Generative-AI analyst support provided contextual summarization and investigation coordination. SOAR playbooks executed governed containment actions. Analyst supervision remained part of the response process, including validation, escalation, and high-impact decision-making.

Several limitations were identified. The evaluation was conducted within a single enterprise-style environment and primarily within one vendor ecosystem, and the exact internal model algorithms were proprietary. Certain social-engineering vectors involving legitimate services and password-protected content were not detected during the initial delivery stage. The false-positive activity comparison was descriptive because the two environments used different classification mechanisms. Behavioral models also required an initial calibration period before stable operational performance was achieved. These observations indicate the continued need for governance procedures, layered controls, and human oversight in AI-supported security operations.

The evaluated architecture is applicable to cloud-dependent environments processing sensitive operational and user-related data across distributed cloud applications. These environments combine continuous security monitoring, incident correlation, automated response, and governance mechanisms to support secure operation.

Future work will consider evaluation in heterogeneous multi-cloud environments, longer observation periods, controlled replay of comparable attack traces, ablation analysis of individual architecture components, cross-vendor comparison of XDR platforms, evaluation of different generative-AI backends, and analysis of adversarial AI techniques and adaptive response mechanisms.

Author Contributions: Conceptualization, A.C. and T.B.; Methodology, A.C.; Validation, F.S., V.S. and A.I. (Andriana Ivanova); Formal analysis, A.I. (Antonina Ivanova), F.S., V.S. and A.I. (Andriana Ivanova); Resources, A.C.; Data curation, A.C.; Writing—original draft, A.C.; Writing—review & editing, T.B., A.I. (Antonina Ivanova), F.S. and V.S.; Visualization, A.I. (Antonina Ivanova) and A.I. (Andriana Ivanova); Supervision, T.B. and A.I. (Antonina Ivanova); Project administration, T.B.; Funding acquisition, F.S. All authors have read and agreed to the published version of the manuscript.

Funding: The article processing charge was funded by the South-West University “Neofit Rilski”, Blagoevgrad, Bulgaria. The research is related to the activities of Project KP-06-N100/7/10.12.2025.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The datasets generated and analyzed during the current study are not publicly available due to security and confidentiality restrictions associated with the evaluated enterprise cloud environment. Aggregated results and methodological information relevant to the study are provided in the manuscript.

Acknowledgments: This research was mainly carried out within the framework of Project KP-06-N100/7/10.12.2025 “Intelligent Adaptive Approaches to Supporting Primary Education for SEN through Machine Learning and Robotic Technologies”. The authors acknowledge the project’s focus

on secure, reliable, and adaptive digital infrastructures supporting educational technologies for learners with special educational needs.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ACR	Automated Closure Rate
EDR	Endpoint Detection and Response
FPA	False-Positive Activity
IaaS	Infrastructure as a Service
ICR	Incident Correlation Ratio
MTC	Mean Time to Closure
MTT	Mean Time to Triage
MTTD	Mean Time to Detect
MTTR	Mean Time to Respond
PaaS	Platform as a Service
SIEM	Security Information and Event Management
SOAR	Security Orchestration, Automation, and Response
SOC	Security Operations Centre
TPR	True-Positive Rate
UEBA	User and Entity Behavior Analytics
VPN	Virtual Private Network
XDR	Extended Detection and Response

References

1. Gartner. Gartner Forecasts Worldwide Public Cloud End-User Spending to Total \$723 Billion in 2025. Available online: <https://www.gartner.com/en/newsroom/press-releases/2024-11-19-gartner-forecasts-worldwide-public-cloud-end-user-spending-to-total-723-billion-dollars-in-2025> (accessed on 10 May 2026).
2. SentinelOne. Cloud Security Statistics 2024. Available online: <https://www.sentinelone.com/cybersecurity-101/cloud-security/cloud-security-statistics/> (accessed on 10 May 2026).
3. Tariq, S.; Chhetri, M.B.; Nepal, S.; Paris, C. Alert Fatigue in Security Operations Centres: Research Challenges and Opportunities. *ACM Comput. Surv.* **2025**, *57*, 224. [CrossRef]
4. Mgbemele, A.F. Advancing Cyber Threat Detection through SIEM-Based Automation and MITRE ATT&CK Aligned Analytics: A Systematic Review. *Asian J. Res. Comput. Sci.* **2026**, *19*, 233–254. [CrossRef]
5. Nasim, S.; Pranav, P.; Dutta, S. A Systematic Literature Review on Intrusion Detection Techniques in Cloud Computing. *Discov. Comput.* **2025**, *28*, 107. [CrossRef]
6. Sowmya, T.; Mary Anita, E.A. A Comprehensive Review of AI Based Intrusion Detection System. *Meas. Sens.* **2023**, *28*, 100827. [CrossRef]
7. ENISA. *ENISA Threat Landscape 2023*; European Union Agency for Cybersecurity: Heraklion, Greece, 2023. Available online: <https://www.enisa.europa.eu/publications/enisa-threat-landscape-2023> (accessed on 10 May 2026).
8. Akamai. New Study Finds 84% of Security Professionals Experienced an API Security Incident. Available online: <https://www.akamai.com/newsroom/press-release/new-study-finds-84-of-security-professionals-experienced-an-api-security-incident-in-the-past-year> (accessed on 10 May 2026).
9. ISC2. ISC² 2024 Cybersecurity Workforce Study. Available online: <https://www.isc2.org/Insights/2024/10/ISC2-2024-Cybersecurity-Workforce-Study> (accessed on 10 May 2026).
10. Ban, T.; Takahashi, T.; Ndichu, S.; Inoue, D. Breaking Alert Fatigue: AI-Assisted SIEM Framework for Effective Incident Response. *Appl. Sci.* **2023**, *13*, 6610. [CrossRef]
11. González-Granadillo, G.; González-Zarzosa, S.; Diaz, R. Security Information and Event Management (SIEM): Analysis, Trends, and Usage in Critical Infrastructures. *Sensors* **2021**, *21*, 4759. [CrossRef] [PubMed]
12. CrowdStrike. AI SIEM: Next-Gen SIEM Powered by AI. Available online: <https://www.crowdstrike.com/en-us/cybersecurity-101/next-gen-siem/aisiem/> (accessed on 10 May 2026).
13. Thandayutham, K. Federated Security Control Data Fabric: Scalable Telemetry Normalization and Orchestration in Multi-Cloud Environments. *Int. J. Cloud Comput. Serv.* **2025**, *11*, 4414. [CrossRef]

14. Palo Alto Networks. What is Extended Detection and Response (XDR)? Available online: <https://www.paloaltonetworks.com/cyberpedia/what-is-extended-detection-response-XDR> (accessed on 10 May 2026).
15. Gartner. Market Guide: Extended Detection and Response. Available online: <https://www.gartner.com/reviews/market/extended-detection-and-response> (accessed on 10 May 2026).
16. Meng, X. Advanced AI and ML Techniques in Cybersecurity: Supervised and Unsupervised Learning, Reinforcement Learning and Neural Networks in Threat Detection and Response. *Appl. Comput. Eng.* **2024**, *82*, 24–28. [[CrossRef](#)]
17. Diogenes, Y.; DiCola, N.; Turpijn, T. *Microsoft Azure Sentinel: Planning and Implementing Microsoft's Cloud-Native SIEM Solution*; Microsoft Press: Redmond, WA, USA, 2022.
18. Pitkar, H. Cloud Security Automation Through Symmetry: Threat Detection and Response. *Symmetry* **2025**, *17*, 859. [[CrossRef](#)]
19. Boughton, D.; Reid, I. The Role of Artificial Intelligence in SOC Operations: Adoption, Perception, and Workforce Impact. In *Proceedings of the 11th International Workshop on Socio-Technical Perspectives in Information Systems (STPIS 2025)*; CEUR Workshop Proceedings: Aachen, Germany, 2025; Volume 4134, p. 28.
20. Winkler, A.M.; Sharma, P. Proactive Threat Detection in Enterprise Systems Using Wazuh: A MITRE ATT&CK Evaluation. *Comput. Secur.* **2025**, *159*, 104702. [[CrossRef](#)]
21. Farzaan, M.A.M.; Ghanem, M.C.; El-Hajjar, A.; Ratnayake, D.N. AI-Enabled System for Efficient and Effective Cyber Incident Detection and Response in Cloud Environments. *arXiv* **2024**, arXiv:2404.05602. [[CrossRef](#)]
22. NIST. *Framework for Improving Critical Infrastructure Cybersecurity (NIST CSF v1.1)*; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2018. [[CrossRef](#)]
23. Rose, S.; Borchert, O.; Mitchell, S.; Connelly, S. Zero trust architecture. *NIST Spec. Publ.* **2020**, *800*, 1–52. [[CrossRef](#)]
24. Lilhore, U.K.; Simaiya, S.; Alroobaea, R.; Baqasah, A.M.; Alsafyani, M.; Alhazmi, A.; Khan, M.M. SmartTrust: A Hybrid Deep Learning Framework for Real-Time Threat Detection in Cloud Environments Using Zero-Trust Architecture. *J. Cloud Comput.* **2025**, *14*, 35. [[CrossRef](#)]
25. Teece, D.J. Explicating Dynamic Capabilities: The Nature and Microfoundations of (Sustainable) Enterprise Performance. *Strateg. Manag. J.* **2007**, *28*, 1319–1350. [[CrossRef](#)]
26. Creswell, J.W.; Plano Clark, V.L. *Designing and Conducting Mixed Methods Research*, 3rd ed.; SAGE Publications: Thousand Oaks, CA, USA, 2017.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.