

Cognitive Cloud Defense: AI Frameworks for Enterprise Threat Intelligence

Sophia Martinez
Asst Professor

Abstract

The rapid evolution of enterprise information technology has increased the complexities of information systems to levels that present augmented challenges for proactive identification and mitigation of security vulnerabilities by information security practitioners and organizations. The growing volume and increasing sophistication of attacks against enterprise infrastructure are exposing vulnerabilities in traditional protection measures. Consequently, enterprise security teams face a rising tide of cyber incidents that, despite major investments in security-prevention measures, continues to exceed their ability to control.

The present study introduces an artificial intelligence–machine-learning modelling framework designed to enable and assist enterprises in the detection and prevention of the majority of malware and APT attack vectors. The modelling framework employs, generates, and uses AI-driven threat intelligence that encompasses extensive visual and searchable data representations of the attack techniques, tactics, and procedures (TTPs) of cyber adversaries across the globe. The modelling framework was evaluated by comparing its detection efficacy and average time taken to detect both traditional malware-based attacks and advanced persistent threats against those of an organization that is globally recognized as a leader in providing security for enterprise infrastructure. The results revealed a notable reduction in average time taken to detect advanced persistent cloud threats over a more conventional and established security control approach.

Keywords: Enterprise Cybersecurity, Artificial Intelligence in Security, Machine Learning for Threat Detection, Malware Detection, Advanced Persistent Threats (APT), Threat Intelligence Platforms, AI-Driven Threat Intelligence, Cyber Attack Detection, Security Vulnerability Management, Proactive Threat Mitigation, Attack Techniques Tactics and Procedures (TTPs), Cyber Threat Analytics, Security Operations Centers (SOC), Intelligent Security Frameworks, Real-Time Threat Detection, Cloud Security Threats, Cyber Defense Automation, Security Incident Response, Predictive Threat Modeling, Enterprise Security Architecture.

1. Introduction

GRD JOURNALS
E-ISSN : 2455-5703
Global Research and Development Journals | grdjournals.com
info@grdjournals.com
Volume-04 Issue-02 APRIL-JUNE 2026
Received: APRIL 20 Revised: MAY 03 Accepted: MAY 25 Published: JUNE 08

Securing data while it is processed, stored, or transmitted in public or private IaaS clouds is extremely challenging. The Cloud Security Alliance lists the top risks to Cloud Computing Security, noting that a majority are correlated with insufficient security controls in the cloud. In many cases, this stems from a basic understanding of the deployment security requirements, security features of key cloud enablers DevOps, SecOps, and SRE, and the existence of public and private tools, patterns and frameworks. Artificial Intelligence (AI) and Machine Learning (ML) frameworks use AI and ML techniques to gather data from multiple sources and correlate that data to provide an insight into advanced security threats, risks and incidents. Topics for 2023 include the definition and advanced persistent threats shaping the cloud threat landscape, the Tactics, Techniques and Procedures (TTPs) being exploited, and an overview of the supply chain and configuration risks that are expected to pose a threat in the near future.

These alerts are not new – AI systems have been in use in the NatSec arena since the deep blue chess project; the US National Security Agency has been using behavioral signature systems since 1986 and has built up significant expertise. It should not be forgotten that cloud services are a commodity and are being targeted as such; the business models of Security as a Service operators and application service providers are being exploited. The AI engines being used today will eventually be commoditised, and as such, the within industry will evolve to use the AI as a tool, rather than being the cyber threat itself. The risks associated with Detect and Respond readiness to advanced persistent threats in cloud environments. The use of Systems Modelling Language to assess and understand detection and response capability gaps. A consolidated synthesis of all advanced persistent threat landscape assessment work, combined with specific cloud insights, will define Detect and Respond readiness suitability and Capability Maturity Model evaluation across cloud environments, thus highlighting the required improvements and investment.

2. Foundations of AI-Driven Threat Intelligence

AI-driven threat intelligence is defined as the integration of machine learning, deep learning, and natural language processing technologies into enterprise threat analyst intelligence. Broadly, the integration of machine learning into data analytics and analysis for enterprise threat assessment is AI-driven enterprise threat intelligence. Machine learning-driven enterprise threat intelligence uses publicly available sources. Processing such data volumes manually using conventional data analytics tools is not possible. Although public data are trustworthy, there is a need for data integrity, data validation, data governance, and proper data management; these principles must be followed to reduce the exposure of machine learning-enabled enterprise threat intelligence systems to threats. Data provenance can help security analysts source trustworthy



data. Using bleeding-edge data and sensors in such systems pushes them ahead of commercially available detection capabilities with service level agreements of less than 15% false positives.

Sources of data for threat actor Tactics, Techniques, and Procedures (TTPs) include open source intelligence from global trusted resources, cybersecurity vendor sources, global news, privately available blog sites, telegram channels, and dark web data forums. Distillation of threat intelligence analysis of such data in real time and near real time supports the advanced detection of APTs, supply chain attacks, targeted ransomware attacks on critical infrastructure, and, more particularly, attacks on cloud environments in less than five hours with reasoning capabilities and dual-language explanation systems to convey reasoning.

Equation 1) Observability-fusion equation

The strongly emphasizes **cross-layer telemetry fusion** across network, applications, cloud infrastructure, and external intelligence.

Let:

- x_1 = network visibility
- x_2 = application visibility
- x_3 = cloud-control-plane visibility
- x_4 = identity visibility
- x_5 = external threat-intelligence visibility

Then a simple observability-fusion score is:

$$OF = \sum_{j=1}^{5} \alpha_j \, x_j$$

with:

$$\sum_{j=1}^{5} \alpha_j = 1$$

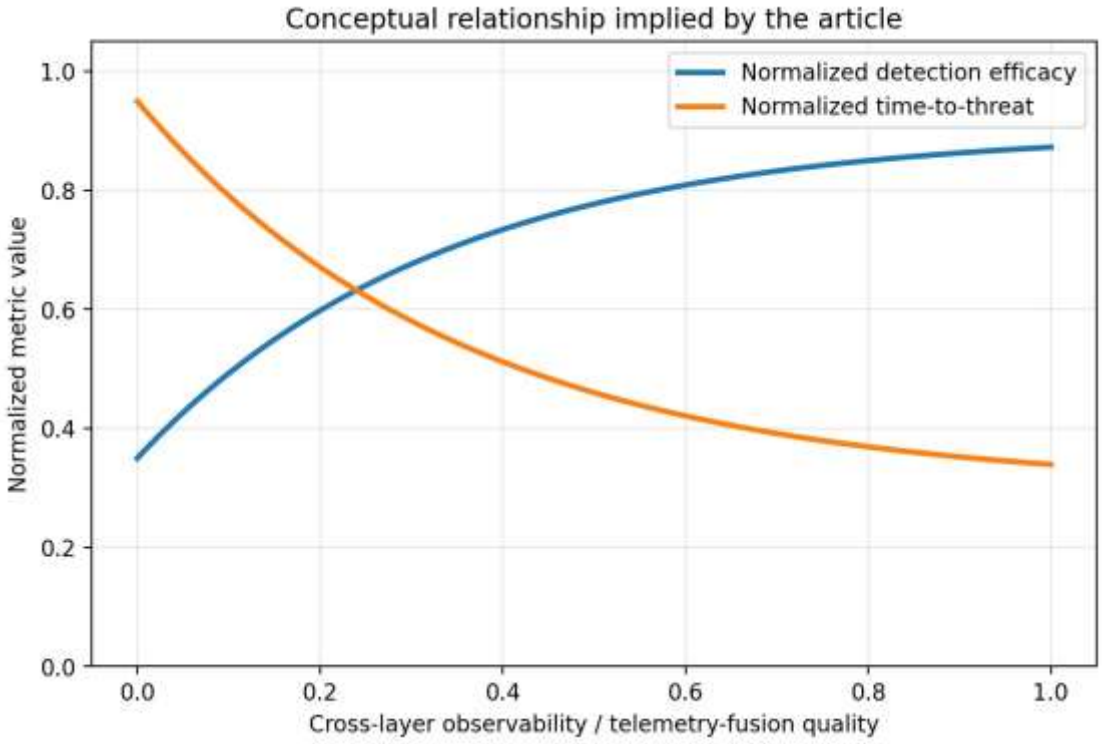
GRD JOURNALS
E-ISSN : 2455-5703
Global Research and Development Journals | grdjournals.com
info@grdjournals.com
Volume-04 Issue-02 APRIL-JUNE 2026
Received: APRIL 20 Revised: MAY 03 Accepted: MAY 25 Published: JUNE 08

If we also want to reward **interaction** between layers, use:

$$OF^* = \sum_{j=1}^{5} \alpha_j \, x_j + \lambda \sum_{p<q} x_p \, x_q$$

The extra term:

$$\lambda \sum_{p<q} x_p \, x_q$$



2.1. Definitions and scope

Scope encompasses data sources, concepts, and analytic objectives within AI-driven threat intelligence. These definitions enable investigation of enterprise AI-enabled cloud security

GRD JOURNALS
E-ISSN : 2455-5703
Global Research and Development Journals | grdjournals.com
info@grdjournals.com
Volume-04 Issue-02 APRIL-JUNE 2026
Received: APRIL 20 Revised: MAY 03 Accepted: MAY 25 Published: JUNE 08

frameworks, especially architectures for data observability. Threat intelligence identifies resources likely to sustain advanced threats, thereby enhancing protection. Advanced threat campaigns demonstrate sophisticated know-how and systematic investigations of specific targets over extended periods. Capturing advanced threat techniques, tactics, and procedures (TTPs) facilitates comprehensive defense strategies. Intelligence expands beyond threat detection to encompass resilience, adaptiveness, and continuity, continuously informing security, risk, and assurance processes. Cybersecurity threat intelligence utilizes diverse sources—external open, proprietary, and commercial—as well as internal indicator repositories, telemetry, and specialized sensors. Data provenance and quality underpin effective intelligence operations, guiding protective and remedial enhancements.

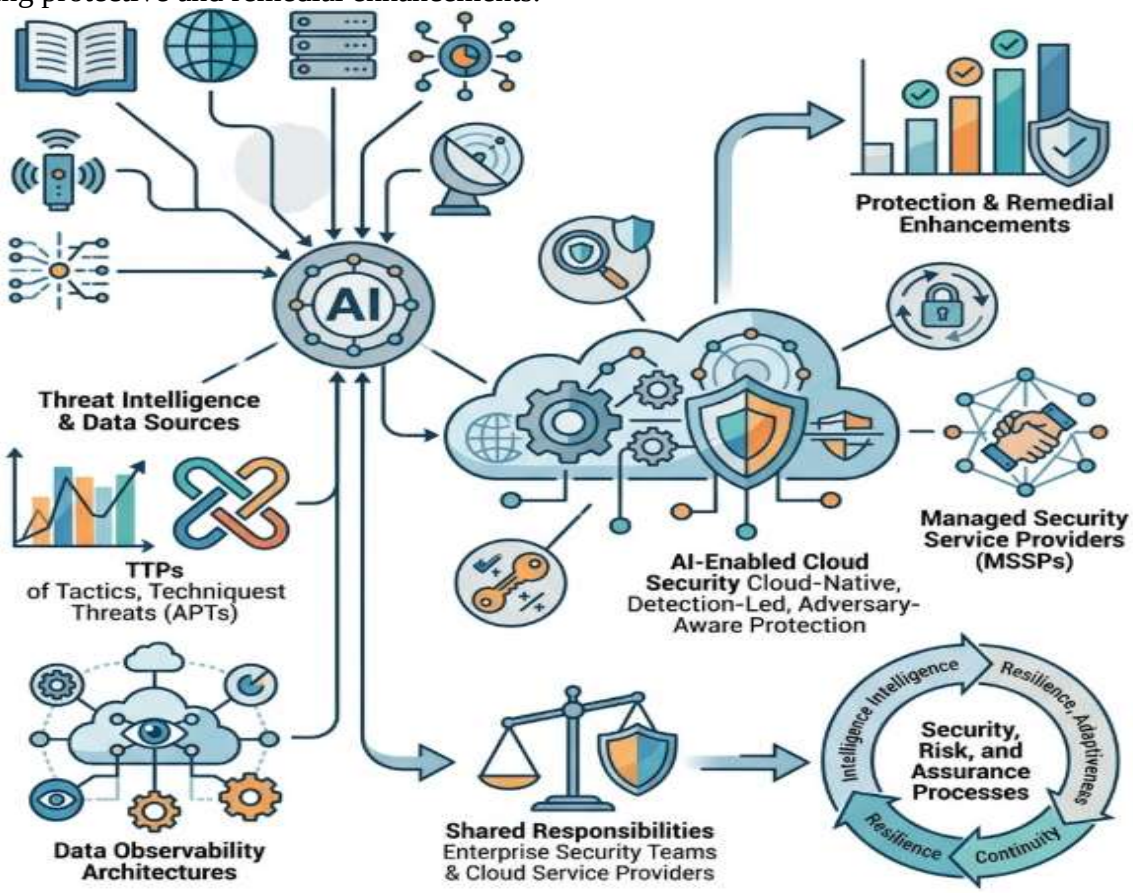




Fig 1: Synergizing AI-Driven Threat Intelligence and Cloud-Native Data Observability: A Framework for Enterprise Advanced Threat Detection, Remediation, and Shared Responsibility

AI-enabled cloud security frameworks represent a growing class of products and solutions delivering cloud-native, detection-led, adversary-aware protection. These technologies operate synergistically with Managed Security Service Provider (MSSP) offerings to discover, identify, contain, and remediate advanced threats—specifically advanced persistent threats (APTs)—present in cloud environments. Shared responsibilities transfer some detection burden from enterprise security teams to cloud service providers, improving time-to-threat and, by extension, reducing risk exposure. To be effective, however, these detection-led services must be purpose-built for public cloud environments, designed to discover advanced threats that invest significant time during pre-attack and attack reconnaissance.

2.2. Key AI techniques in threat intelligence

AI-driven threat intelligence leverages a range of AI techniques and tools, including supervised, unsupervised, and reinforcement learning, as well as computer vision, natural language understanding, and semantic analysis; data mining; big-data analytics; sophisticated graphical and streaming-query-processing systems; replay and event-correlators; and rule-based engines. Data sources span structured, semi-structured, and unstructured data, and natural-language-based information feeds, covering both operational data (e.g., attack indicators and alerts) and information pertinent to the planning, preparation, and execution by threat actors. Analytic goals focus on identifying trusted and authoritative data sources; detecting, categorizing, and predicting threat actors, their capabilities, and the techniques, tactics, and procedures; and selecting and synthesizing relevant and reputable information on actors and associated TTPs.

Analyses combine data sourced from multiple domains and functional silos; define an integrated architecture that enriches the deep storage, consolidated views, and consistency of all data relevant to ICY (what is happening now), POA (what might happen next), and TLR (what has happened); span security, risk, compliance, privacy, and fraud reduction; and increase visibility into activities both above and below the organizational boundary. The resultant end-to-end observability and contextual visibility enhance detection performance, automating detection and response for known threat events and significantly accelerating time-to-threat for unknown threats. These decisions inform the modeling and simulation of attacks against known vulnerabilities, misconfigurations, and dependencies.



3. Methodology

The overall research design employs a multiphase approach. The first phase comprises a meta-analysis synthesizing the detection and prevention capabilities of AI-enabled frameworks deployed to address security risks across cloud delivery models. For the second phase, the temporal emphasis shifts to explore the detection and containment capabilities being integrated into enterprise security operations for threat pathways across hybrid cloud ecosystems.

The analysis of AI-enabled detection and prevention capabilities across the cloud delivery models uses a meta-analysis methodology. Detection, prevention, and observability mechanisms from the reviewed frameworks align with the Major Security Detection Systems structure and are correlated against the threat landscape evolving around the emergence of cloud-native architectures to ascertain their overall efficacy. The results of the analysis are adapted and scaled to articulate a broader view of detection capabilities for enterprise security operations.

The VDI model for risk detection and mitigation performance represents an application-centric view of incidents and threats that categorizes them along a continuum defined by increasing visibility and complexity. The synthesis approach correlates core VDI detection and risk containment practices across the cloud-native operating model and AI-enabled observability of hybrid multicloud estates against the evolving APT risk trajectory with a particular focus on enabling enterprise detection, disruption, and defence in depth.

3.1. Framework for Data Oversight and Provenance

Data sourcing, trustworthiness, lineage, integrity, traceability, and appropriateness are foundational concerns for any observable or understandable system. Data governance architecture therefore encompasses the data itself, the processes supporting collection, storage, and sharing, and business services guiding the journey from design to data consumption and modelling. Control over data is critical for downstream decision-making and regulatory compliance; for AI models, proper supervision of the training, validation, and evaluation datasets is necessary to avoid issues such as bias, integrity, or provenance gaps.

Cloud-native services introduce specific challenges for data oversight, notably because data commonly reside in third-party environments or under the control of cloud service providers (CSPs). Clouds also introduce new signals and telemetry streams, highlighting the need for responsible governance along all axes and throughout the data lifecycle. Consequently, cloud-native observability supports higher-quality models deployed in lower-latency fashions, yet

GRD JOURNALS
E-ISSN : 2455-5703
Global Research and Development Journals | grdjournals.com
info@grdjournals.com
Volume-04 Issue-02 APRIL-JUNE 2026
Received: APRIL 20 Revised: MAY 03 Accepted: MAY 25 Published: JUNE 08

achieving that goal requires attention to the union of signalling across layers and the shared lineage of observability solutions and consolidated data stores. AI model performance is ultimately bounded by data: its quality, the appropriateness of the information fed into the model, and the inclusion of information sources that minimise false positives and enhance detection readiness.

Topic	Final derived equation
Total events	$N = TP + TN + FP + FN$
Detection efficacy / recall	$R = \frac{TP}{TP + FN}$
Precision	$P = \frac{TP}{TP + FP}$
False positive rate	$FPR = \frac{FP}{FP + TN}$
False negative rate	$FNR = \frac{FN}{TP + FN} = 1 - R$
Specificity	$\frac{TN}{TN + FP} = 1 - FPR$
F_1-score	$F_1 = \frac{2PR}{P + R} = \frac{2TP}{2TP + FP + FN}$
Per-attack time-to-threat	$TTT_i = t_i^{detect} - t_i^{start}$
Mean time-to-threat	$\overline{TTT} = \frac{1}{k}\sum_{i=1}^{k}\left(t_i^{detect} - t_i^{start}\right)$
Relative TTT improvement	$\frac{\overline{TTT}_{baseline} - \overline{TTT}_{AI}}{\overline{TTT}_{baseline}} \times 100$
Alert burden	$A = TP + FP$
Analyst workload	$W = \tau(TP + FP)$
Wasted analyst fraction	$\frac{FP}{TP + FP}$

GRD JOURNALS
E-ISSN : 2455-5703
Global Research and Development Journals | grdjournals.com
info@grdjournals.com
Volume-04 Issue-02 APRIL-JUNE 2026
Received: APRIL 20 Revised: MAY 03 Accepted: MAY 25 Published: JUNE 08

Topic	Final derived equation
Data quality score	$DQ = w_1T + w_2L + w_3I + w_4V + w_5G$
Geometric provenance score	$DQ_g = T^{w_1}L^{w_2}I^{w_3}V^{w_4}G^{w_5}$
Observability-fusion score	$OF = \sum_{j=1}^{5} \alpha_j x_j$
Synergistic observability	$OF^* = \sum \alpha_j x_j + \lambda \sum_{p<q} x_p x_q$
Composite effectiveness	$S = \beta_1 DQ + \beta_2 OF + \beta_3 R + \beta_4 P - \beta_5 FPR - \beta_6 Drift$
Drift	$Drift = d\left(P_t(x), P_{t+\Delta}(x)\right)$
Mean-shift drift	(Drift_\mu=\frac{1}{m}\sum_{j=1}^{m}
Risk score	$Risk = \gamma_1 Likelihood + \gamma_2 Impact + \gamma_3 Exposure - \gamma_4 ControlStrength$
Human-in-the-loop confidence	$p_{final} = \eta p_{model} + (1-\eta)h$

4. Objective of the Study

The objective is to assess the impact of AI and ML technologies on enterprise security postures in the context of cloud computing for both data storage and critical infrastructure. The aim is twofold: First, to evaluate the current state of enterprise cloud security. Second, to align the evaluation with the threat landscape and identify changes and emerging risks. The analysis builds on the threat assessment included in the Cloud Threat Landscape Report 2023 by Crowdstrike The enterprise risk and security landscape continues to evolve. Incidents are becoming more severe, sophisticated, targeted, and integrated, and operating in cloud-native, multi-cloud, and hybrid-deployed environments. Cybercriminals are engaged in business and share information about successful attacks to scale operations. Todays cybercriminals have more risk and are pushing the boundaries of crime to generate more impact on their targets. APPTs have formed as a group effort to succeed in more complex IOCs and TTPs.



The hypotheses support this assessment: AI and ML implementations are reducing enterprise detection time for TTPs addressed by the detection solutions, yet they remain a small percentage of the overall enterprise security footprint. They do not currently address the total enterprise risk landscape, nor do they account for other areas and advanced technical controls but support detection in those areas. Such risk changes provide guidance on where enterprise defenses should focus next, dictating investments in and the topology for supporting AI/ML in the enterprise risk and security posture.

Equation 2) Data governance / provenance equation derived from the architecture

The repeatedly stresses that AI model quality depends on data trustworthiness, lineage, integrity, validation, and provenance.

A direct formalization is a weighted quality score:

$$DQ = w_1T + w_2L + w_3I + w_4V + w_5G$$

where:

- T = trustworthiness
- L = lineage traceability
- I = integrity
- V = validation quality
- G = governance compliance

with:

$$w_1 + w_2 + w_3 + w_4 + w_5 = 1$$

Why this form?

Because the treats provenance as a multi-factor requirement, not a single binary feature.

If all components are normalized in $[0,1]$, then:

$$0 \leq DQ \leq 1$$



1.1 Geometric version for stricter governance

If one weak factor should sharply penalize the score, use geometric aggregation:

$$DQ_g = T^{w_1} L^{w_2} I^{w_3} V^{w_4} G^{w_5}$$

4.1. Purpose and Goals of the Research Study

The objective of this research study is to provide a comprehensive collection of AI-enabled cloud security frameworks that would address the most severe cloud-focused security threats presently affecting enterprises. This effort is bolstered by an analysis that examines architectural selection principles and current limitations. Threat actors continuously evolve their techniques, tactics, and procedures, resulting in the emergence of new threats and shifting cloud security requirements. The following four research questions provide the guiding calculations and objectives. The first question investigates the specific AI-enabled cloud security frameworks implemented by major enterprises. The second addresses the role that advanced persistent threat actors play when exploiting cloud environments for sustained cyber-espionage operations. The third question examines the vulnerabilities associated with supply chain and configuration issues in the cloud. Finally, the fourth evaluates the efficacy of a detection framework by recording false-positive rates and measuring the time taken to detect actual malicious activity.

This research makes a valuable contribution to information security monitoring in general and cloud security monitoring in particular. A consolidated summary of significant AI-enabled cloud security frameworks, a threat-landscape analysis focusing on advanced persistent threats, supply-chain vulnerabilities, potential sticking points in cloud-native deployments, and detection-performance-benchmarking data for a leading detection framework are now available. An architectural review highlighting these areas enables security practitioners to make better-informed and more-conscious decisions. Threat actors are invariably setting the pace for risk and damage, and organisations globally are at the hopeless end of a catch-up game.

5. Research Summary

Data policies and governance structures remain largely undeveloped, and a pessimistic attitude toward data-sharing initiatives endures, especially as data volume and sensitivity increase. Advanced persistent threat actors have adapted their tactics, techniques, and procedures to take advantage of cloud technologies and platforms but have yet to be systematically profiled in publicly available resources. Supply chain and configuration risks will remain a major source



of attacker opportunity, as dependencies grow deeper, B-grade libraries are used at scale, and organisations struggle to understand the scope of risk represented by complex and fragile infrastructure-as-code deployments. The tension between business need, security operation capability, and regulatory obligation will push noise-reduction progress to the fore. As businesses continue shifting to a cloud-first mentality, extensive resources will be deployed to tap into the underlying technology's responsiveness and scale.

Recommendation for large-scale deployments of native cloud workloads and use of hybrid infrastructures has been driven by research compiled from interviews with major enterprises, implementation partners, and platform vendors. Key focus areas—training, platform service governance, and advanced telemetry integration—have been identified as vital to maintaining visibility and risk management within hybrid and multicloud environments. Increasingly mature platforms offer enhanced risk-service, data-privacy, and regulation-alignment features, and successful implementation requires targeted resourcing across these categories. Requirements and best practices for accurately monitoring an insecure-by-design platform-as-a-service offering are poorly understood, and the processes necessary to maintain these capabilities remain largely undocumented within security-operations centres.

5.1. Comprehensive Analysis of AI-Driven Cloud Security Architectures

The evolving cloud landscape introduces myriad complexities for security. While many security domains have made significant advances owing to AI, the research community still confronts the challenge of stitching together cross-layer telemetry signals from multiple sources in a single unified solution. Current solutions have started to take advantage of these cross-layer signals, however, they appear limited in terms of detection efficacy or time-to-threat metrics. The risk of undetected advanced persistent threats in cloud environments is rising rapidly as APT actors formulate additional attack tactics, techniques, and procedures specific for cloud environments.

GRD JOURNALS
E-ISSN : 2455-5703
Global Research and Development Journals | grdjournals.com
info@grdjournals.com
Volume-04 Issue-02 APRIL-JUNE 2026
Received: APRIL 20 Revised: MAY 03 Accepted: MAY 25 Published: JUNE 08

These disrupting TTPs demand special focus on detection.

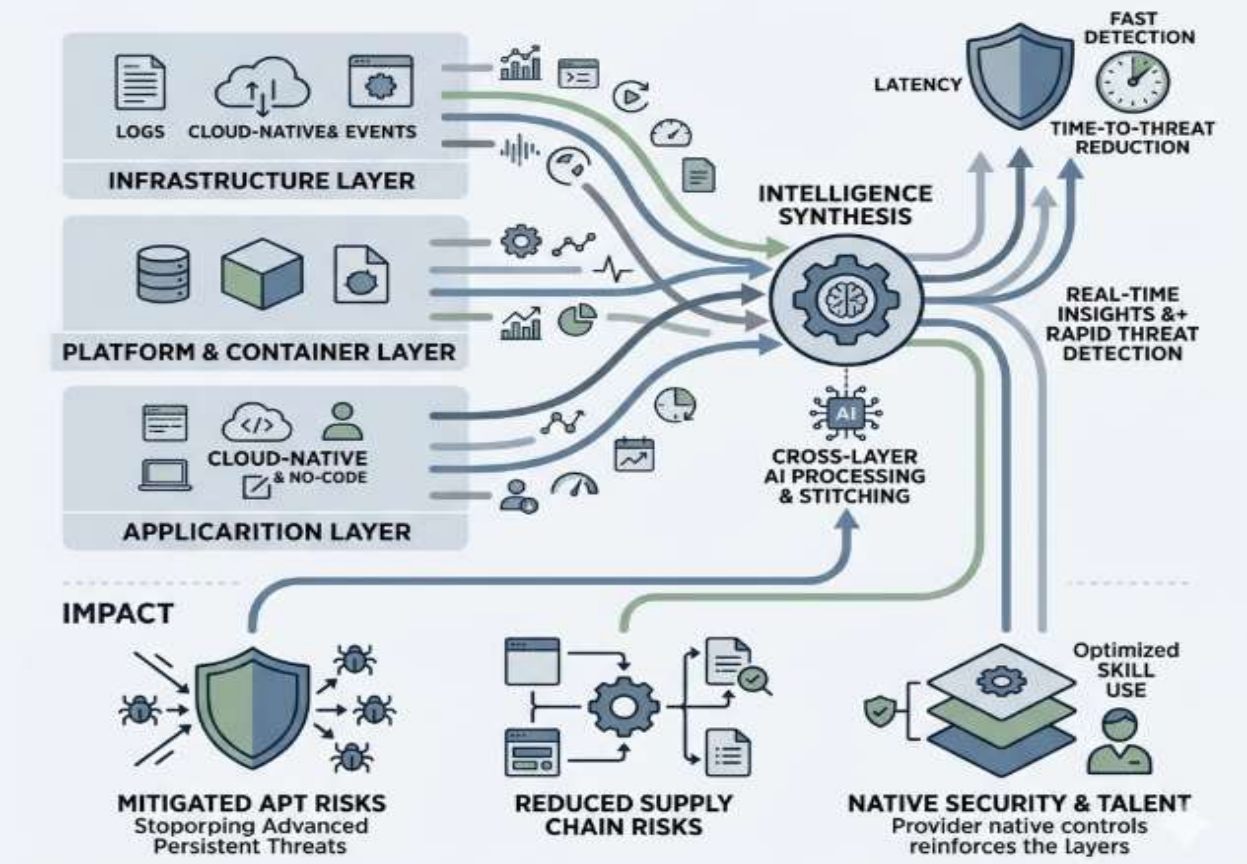


Fig 2: Bridging the Telemetry Gap: AI-Driven Cross-Layer Synthesis for Mitigating Cloud-Native APTs and Supply Chain Vulnerabilities

The soaring popularity of low/no-code development frameworks, combined with ever-increasing pressures to optimize costs, is driving a swift move toward cloud-native application development. Such development offers better velocity and cost efficiency, however, introduces its own set of risks—from poor supply chain practices to inherent misconfigurations. As organizations take advantage of multiple clouds, the associated talent gap has led security to become an afterthought within numerous cloud-native applications. Nonetheless, the providers are starting to provide native security features—aided by machine learning and artificial



intelligence—to bolster these environments, with some leading organizations testing the waters to prove out these security controls.

6. Architectural Review of AI-Enabled Cloud Security Frameworks

1. Data governance and provenance

Assessment of data sourcing, trust, lineage, and integrity across cloud contexts.

Operationalization of European Union's General Data Protection Regulation (GDPR) and associated privacy regulations is nearing execution, positioning enterprise governance as a top priority for both consumers and regulators. Customer-scale and generalized AI-enhanced applications depend heavily on data ingestion, regulation compliance, sufficient analytical performance, and user expectations. Organizations have either dedicated strategies or summonsed board attention in these areas, seeking best practice and appetite-dependent trade-offs and presenting risk and reward profiles associated with bias, quality, confidence, label sources, regulation compliance, and operational overhead. Microsoft's chief compliance officer (CCO) highlighted in August 2022 that AI-enhanced services would undergo Burden of Care assessments for highly sensitive applications, as informed by the potential associated consequences. Product reviews raised concerns regarding the implications for chatbots and automated solutions, in citation, warranty, consent, impersonation, public jiggle, and manipulation domains.

Cloud application framework builders are visibly maturing operational accountability and presenting a more perfect answer set concerning data oversight and provenance aspects of public clouds than common-ground clouds. Nevertheless, organizations source a complex web of open-source components, AI-enhanced capabilities, and composite-as-a-service architectures, with no common oversight and the internecine protection whispers. Platform-ai-enhanced services remain favoured in a risk-versus-reward evaluation by the cloud, although cloud-supplied service-specific spear-phishing bows saw cloud hyperscalers' exposures exploited by human-operated threats and insider actors.

2. Observability and telemetry fusion

Description of integration of signals, telemetry streams, and cross-layer visibility.



Wide visibility across application stacks, from end-user experience to infrastructure health, is an elemental compressive foundation now foundational to remediation efforts. Enterprises affect vendors' success or demise, necessitating initial perspectives. Cross-signal stream extraction of any observability platform from wholly public domain consumable or composite libraries orchestrating closet-white-label codebases plays a similar role in remediation efforts, replacing the rescan execution of the half-life threat surface in detection and response. Fortune 100 and non-financial Global 50 enterprises display publicly exposed components in exposures, echoing attacks on mobile platform enterprises exploiting high-visibility consumer applications as the attack surface relocates to external-facing assets. Integration of Hypertext Transfer Protocol Secure (HTTPS) Web Application Federation with Transport Layer Security (TLS) replaces the concept of redirection between different web applications owned by different parties, without the overhead of encrypted data tokenisation, while RiskIP consolidates scanning of the installation air gap.

6.1. Data governance and provenance

The Governance, Risk, and Compliance (GRC) domain has generated a plethora of requirements and best practice frameworks addressing various verticals, industries, and regulatory mandates. However, when it comes to AIOps platforms, particularly those vitally important for an organization's security posture, many enterprises prioritize the deployment of such systems without sufficient consideration of adequate GRC-oriented data governance foundations. As a result, these AIOps platforms often serve as a collection of point solutions driven primarily by their respective vendors' proprietary data sources. Consequently, issues related to data sourcing, trustworthiness, lineage, and integrity within the GRC domain remain poorly understood, with little consideration for how such concerns might utilize AIOps frameworks in public, private, and hybrid cloud deployments.

Taking into account the observed reputation and credibility of the data sources, an information-centric perspective to security and risk analysis indicates the lack of proper data governance principles in many AIOps deployments. As data provenance principles are increasingly becoming an essential requirement for regulatory compliance, it comes as no surprise that the lack of AIOps data governance is giving rise to major concerns, particularly for those use cases that utilize patently untrusted data sources, such as in the context of enterprise threat intelligence vendors and their users.

6.2. Observability and telemetry fusion



The integration of novel observability signals and telemetry sources achieved such richness that orthogonal streams from networks, applications, and cloud infrastructure could sew together a sensible view of activities surviving deeper inspection. However, advanced telemetry datasets still required considerable effort for coherent fusion. Several directions were pursued for better orchestrated primary collections, enhanced cross-layer visibility, and seamless enrichment via signature or analytics-based correlation. Mechanisms available to integrate dependency-based models supported visibility over the orchestration layer and service mesh underpinnings, with important implications both for attack detection and remediation ventures. Meanwhile, traffic-stream and application-layer protocols remained largely decoupled, in spite of the enrichment power that complete traffic visibility might bring and the evolving trends towards codeless or artificial path creation.

Telemetry steering still depended on a sacrificed-third-party vision, with important opportunities left on the table for skipping such external constraints and probing the internal organization of application-layer flows. Their increased prevalence over custom encoding and the control they impose on their own infrastructure were argued to warrant an augmented focus on managing their detection-capture elements without violation—adverse conditions being dealt with through managed care and trapping models. Signal overload remained a practical concern. Anti-abuse mechanisms controlled spurious signalling to alert and lift suspicious assets during their passage to the light of reasons, but the ephemeral nature of some miscreants called for stronger containment than mundane alerts. Models and practices borrowed from false-positive mitigation in malware and APT detection surfaced as areas rewarding further consolidation, with external telemetry possibly assisting internal warning management under preservation of the monitoring functions' avoided-third-party character when desired.

Equation 3) Core detection equations derived step by step

The explicitly emphasizes **detection efficacy**, **false positive management**, and **time-to-threat** as the main evaluation dimensions.

3.1 Total-event decomposition

Every observed event is either malicious or benign:

$$N = M + B$$

Every malicious event is either detected or missed:

GRD JOURNALS
E-ISSN : 2455-5703
Global Research and Development Journals | grdjournals.com
info@grdjournals.com
Volume-04 Issue-02 APRIL-JUNE 2026
Received: APRIL 20 Revised: MAY 03 Accepted: MAY 25 Published: JUNE 08

$$M = TP + FN$$

Every benign event is either correctly ignored or wrongly alerted:

$$B = TN + FP$$

Therefore the total observed events can also be written as:

$$N = TP + TN + FP + FN$$

3.2 Detection efficacy

The "detection efficacy" as a central metric. In threat detection literature, that corresponds to the fraction of actual malicious events that are correctly detected. So:

$$\text{Detection Efficacy} = \frac{\text{Detected malicious events}}{\text{All malicious events}}$$

Substitute symbols:

$$R = \frac{TP}{M}$$

Since $M = TP + FN$,

$$R = \frac{TP}{TP + FN}$$

This is also called **Recall** or **True Positive Rate**.

Step-by-step derivation

1. Start with all malicious events: M
2. Split them into detected and missed: $M = TP + FN$
3. Detection efficacy is the detected part over the whole malicious set:

$$R = \frac{TP}{TP + FN}$$



7. Threat Landscape for 2023 and Beyond

Advanced persistent threats in cloud environments

Attackers continue to develop and enhance sophisticated capabilities for targeting cloud environments, and APT tags are now being applied to campaigns in these contexts. Despite the provisioning of security controls and the presence of additional telemetry, many organizations remain reliant on detecting and thwarting threats post-compromise. Microsoft describes the use of deep identity management, continuous and behavioral audit logging, and health monitoring in relation to the company's defense posture against APTs.

Supply chain and configuration risks

The cloud supply chain exposes enterprises to a multitude of external risks that can affect the stability and security of existing services. The speed of development in this market means some relationships can easily become unhealthy and introduce risk into an organization without the organization being aware. Common misconfigurations and errors made by organizations when deploying their applications remain a primary entry vector for attackers and provide opportunities for resource abuse or information theft. Configuring and managing dependencies warrants additional attention, as these dependencies might also introduce risk into an organization.

7.1. Advanced persistent threats in cloud environments

Advanced persistent threat (APT) groups are motivated and technically skilled actors who execute resource-intensive, focused and stealthy attacks for political or economic purposes. The development of cloud environments offers these actors a new platform, with resources that can be exploited for coordinating assaults on third-party organizations or providing the safety of obfuscation and endurance. Research and recording on such activity remains scarce, even as it evolves. While APTs have predominantly targeted specific organizations or national/ethnic groups, a shift toward crime is now being seen, with the aim of generating funding that would enable local geopolitical threats. The scope, scale, and development of these threats, their procedures, techniques, and patterns of behavior, are broadening, driven not only by innovation and the advent of new tools but also by the accessibility of these offerings in underground marketplaces. Consequently, it is essential that the security community take a closer look at this category of cyber criminals, particularly the main drivers of its crime, the underground economy.

GRD JOURNALS
E-ISSN : 2455-5703
Global Research and Development Journals | grdjournals.com
info@grdjournals.com
Volume-04 Issue-02 APRIL-JUNE 2026
Received: APRIL 20 Revised: MAY 03 Accepted: MAY 25 Published: JUNE 08

Moreover, APTs are also increasingly utilizing novel exploit techniques that are sometimes not captured by existing detection mechanisms. While prior studies have principally focused on key issues such as the TTPs of APT groups, the set of skills and duration of exploitation not only remain difficult but also often involve multiple architectures across layers that encompass IoT devices. Such circumstances make APTs rarely seen, but when discovered, their capabilities and TTPs can be exceptionally toxic and alarming. In addition, the ever-changing evolving TTPs of APT groups mean that the security community has to keep up to date with these actors, and complimentary information between different behavioral sources is indispensable for effective alerting. For this reason, the latest report from Group-IB has examined the evidence and skill, bringing together the ecosystem supporting these actors and the detection of cloud resource abuse.

7.2. Supply chain and configuration risks

The managed service model has emerged as a favourite for enterprise outsourcing, enabling cloud providers to build and maintain complex services while consumers rely on preset levels of maturity for basic security hygiene and compliance frameworks. However, widespread reliance on mismanagement of third-party service configurations combined with the existing risk of integrated attack surfaces has given rise to supply chain risks in the highly dynamic cloud ecosystem. New applications are never developed from scratch and almost always rely on a shared component/assembly for the application code (e.g. open-source libraries) and the provisioned platform (e.g. container images, IaC templates, etc.). The growing dependencies on commercial off-the-shelf software (COTS) for value delegations and bundling of solutions has caused independent software vendors (ISVs) to tweak the platform code instead of building on a modular approach, widening the security gap and increasing the TTP overlap between nation-state APTs and cybercrime groups. Detection of vulnerabilities, bad practices, and misconfigurations is more critical than ever in the cloud world. Malicious actors have been observed scanning public repositories like Docker Hub regularly for vulnerable images before launching highly automated attacks against the exploited misconfiguration or dependency.

Automated scanning of configuration and dependency manifests should be used in mission-critical cloud workloads. However, to reduce the number of false positives, it is essential to properly tune these scanners to suit the environment in question, as most tools offer a wide array of sensitive checks that can lead to an avalanche of daily alerts. Exploit systems, such as an evil repository space (called Simply-SNS-less) target a single vulnerability, managing to leverage its source across the enterprise universe at attack speed. Old news on Expo and Google Cloud illustrate how these incidents repeat themselves for any dependency from any corner of



the attack continuum that offers even the tiniest exploit. Exposed unchanged testing accounts, hard-coded secrets, and unexplained TLS issues should also trigger an alert. Detection of cloud storage leaks should occur in theft context (usually via minted credentials) and should hence look for anomalously high volumes of outgoing traffic through storage and sharing services.

8. Case Studies of Enterprise Implementations

The first case study centers on a large-scale deployment in a fully cloud-native environment for a leading global contractor, covering the design, security measures, and outcomes. Alert data from multiple security functions employed within the environment is captured for subsequent analysis using an advanced detection engine augmented with AI-based automated analytics and machine learning, providing an overall detection capability with low false positive rates and reduced mean time to contain detection. The second case study considers a hybrid cloud scenario for a national healthcare system using multi-cloud services. Threat modelling identified the configuration of application access controls and segmenting of the health system's digital properties as key risk areas in a recent penetration test. Controls deployed for detection, response, and specific threat safety are described, together with their use to mitigate a vulnerability identified at the time of the penetration test.

As large-scale cloud-native environments become the norm, providing services at very low cost, the need for increased levels of detection and response system automation becomes paramount to ensure sufficient coverage at a reasonable TCO. …A recent analysis of an initial set of detections and alerts indicated potential for machine learning to automate considerable parts of the intelligence cycle, leading to a deployment model split into multiple centres of gravity with federated learning shared across clouds.

8.1. Large-scale cloud-native deployments

AI-Compose a formal, evidence-based, and objective scholarly narrative that adheres to an academic tone and structured presentation, emphasizing clarity and rigorous argumentation. These capabilities enable deep integration of enterprise security monitoring and detection systems into the cloud provider’s architectures and services, yielding advanced threat detection

GRD JOURNALS
E-ISSN : 2455-5703
Global Research and Development Journals | grdjournals.com
info@grdjournals.com
Volume-04 Issue-02 APRIL-JUNE 2026
Received: APRIL 20 Revised: MAY 03 Accepted: MAY 25 Published: JUNE 08

and response speed, consistency, and scaling effectiveness.

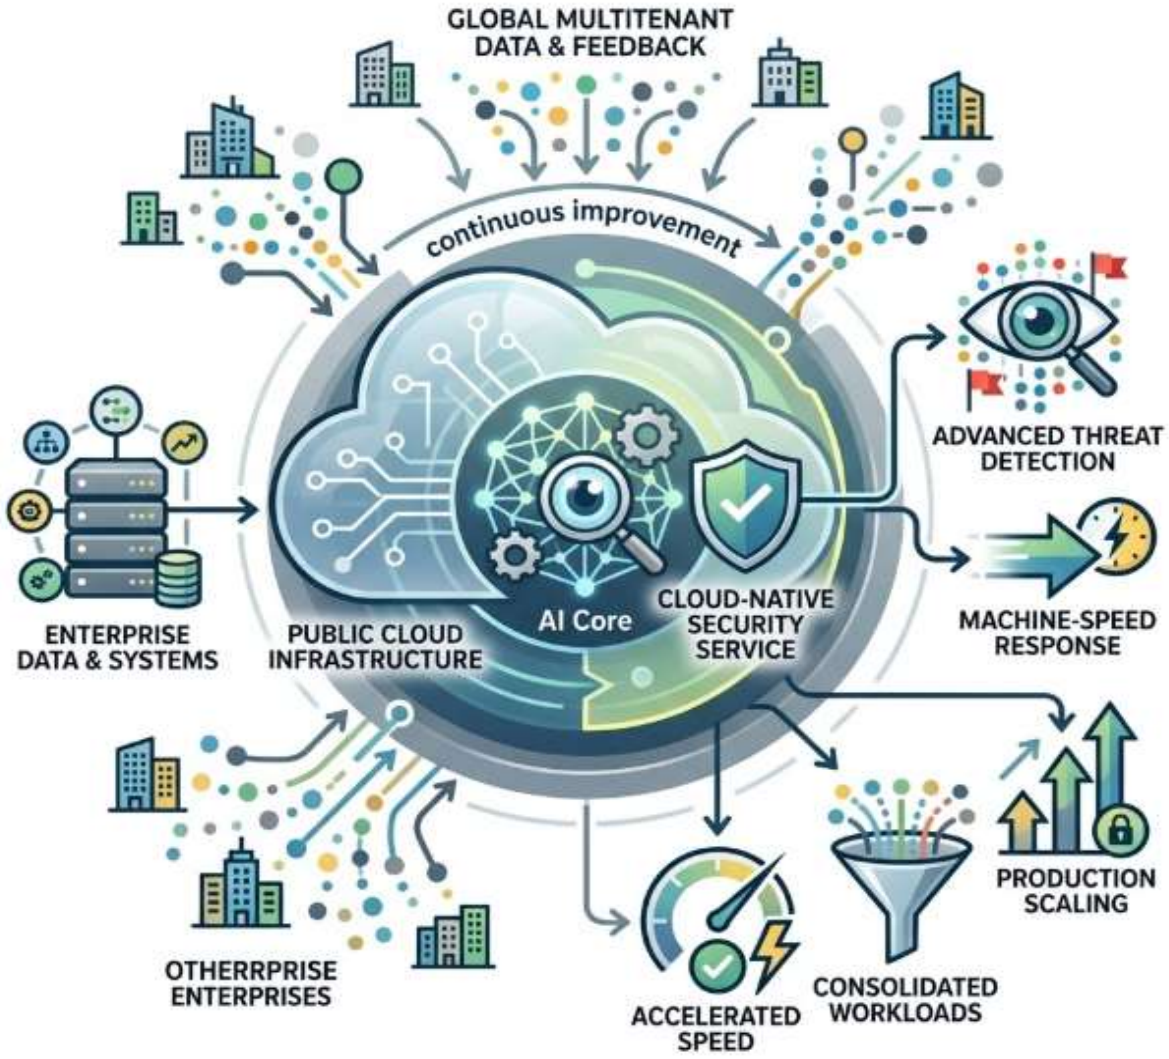


Fig 3: Leveraging Cloud-Native AI for Resilient Enterprise Security: A Comparative Analysis of Integrated Threat Detection and Machine-Speed Response

Public cloud infrastructure represents both a challenge for protecting enterprise data and systems and an opportunity to transform detection and response capabilities. The provider's

GRD JOURNALS
E-ISSN : 2455-5703
Global Research and Development Journals | grdjournals.com
info@grdjournals.com
Volume-04 Issue-02 APRIL-JUNE 2026
Received: APRIL 20 Revised: MAY 03 Accepted: MAY 25 Published: JUNE 08

multitenant environment supplies data from all layers of the infrastructure and responses from many enterprise customers to continually improve detection accuracies while minimizing false positives. These combined benefits can significantly shorten time-to-threats and consolidate detection workloads. The delivery of security monitoring and detection capabilities as services within the public cloud product offer enables enterprise customers to achieve those advantages without deploying additional security monitoring Infrastructure and personnel at enterprise scale.

Three recent enterprise deployments of these capabilities exemplify the accelerated detection and response advantages possible in a pure cloud-native context. The cloud-delivered capabilities have proven effective for threat detection at production scale with efficacy on par with state-of-the-art dedicated enterprise solutions, including the capacity to disrupt true threats at machine speed.

8.2. Hybrid and multi-cloud contexts

Jointly supporting both internal enterprise needs and external customer demands requires a hybrid cloud strategy that leverages orchestration and automation for security and risk management across semitrusted regional and intercontinental partners. The security requirements to support such hybrid cloud strategies have led to high-profile collaborations on cloud security posture management and Cloud-Native Application Protection Platforms; however, despite the documented enterprise-scale cyber appliance implementations, most of the literature continues to focus on either cloud-native or cloud-agnostic designs. Private Cloud-as-a-Service solutions provide another design approach to meet specific regulatory needs and provide additional trust, albeit without the broad capability and geographic footprint of public offerings. These solutions allow enterprise-sensitive workloads to be hosted on a semitrusted HPC infrastructure while the remaining environment and dedicated SLA appliances can use the full-fledged enterprise public cloud. Such a segregation of workloads reduces the attack surface of Cloud Cyber Control and allows the tenants to encrypt data prior to the memory. Security, privacy, sovereignty, and regulatory compliance concerns can be further mitigated by keeping Critical Infrastructure Components on-premise and deploying Security & Management tools within the Critical Domain.

Natural Language Processing plays an essential role in modern threat analytics in hybrid and multi-cloud deployments. Threat detection using Natural Language Processing examines a specialized linguistic description of threat information to improve coverage and reduce false positive messages. Generative AI can create synthetic traffic to build a strong dataset for anomaly detection. It has matured enough that organizations can leverage its power, for example,



in engineering tasks like reverse engineering a threat actor's objective. Natural Language Processing models can categorize the high-velocity alerts generated by devices in a hybrid cloud, while another model classifies the legacy malware family of the command and control Traffic Log. Further analysis showed required modifications to the architecture that would allow them to better address advanced persistent threats.

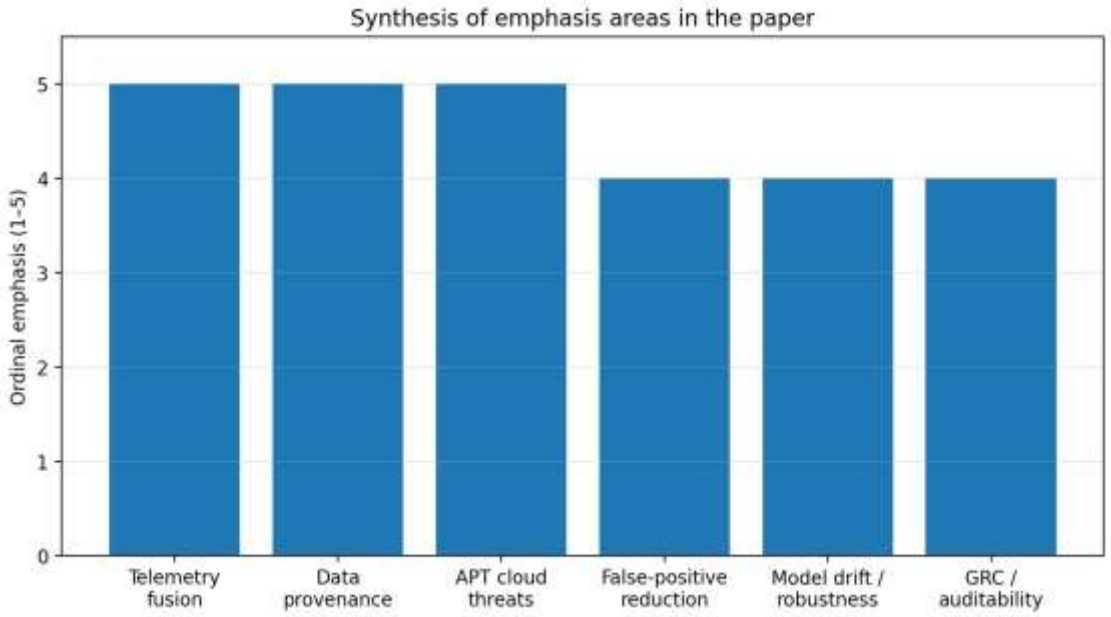
9. Evaluation Metrics and Benchmarking

Detection efficacy and time-to-threat

A comprehensive methodology for evaluating detection performance in cloud security frameworks is presented. The detection metrics and benchmarks are based on established best practices in network security. Deceptive layers within the architecture allow adversarial intelligence and other techniques to inject threats, while an automated evaluation process captures detection events and threat classification to assess performance. Standard evaluation metrics such as detection-rate and time-to-threat, along with scenario-specific benchmarks for non-permanent threats, provide insights into detecting responsiveness and efficacy.

Detection efficacy is of paramount importance in real-time environments like cloud Infrastructure. The abstractions that simplify provisioning and scaling also reduce detection response and mitigation time; thus, detection must encompass not just accuracy but also responsiveness. Baseline detection efficacy is evaluated against established best-practice datasets and metrics, along with technology-agnostic non-permanent baseline scenarios such as Distributed-Denial-of-Service (DDoS) attacks. These standard deviations are essential for scalable deployment-as-a-service within large cloud security frameworks. Secure Intelligence-as-a-Service also addresses hybrid hosting, incorporating edge security by aligning effectiveness with edge-computing best practices.

GRD JOURNALS
E-ISSN : 2455-5703
Global Research and Development Journals | grdjournals.com
info@grdjournals.com
Volume-04 Issue-02 APRIL-JUNE 2026
Received: APRIL 20 Revised: MAY 03 Accepted: MAY 25 Published: JUNE 08



9.1. Detection efficacy and time-to-threat

An expanded set of detection-related metrics is needed for efficiency and time-to-threat that intersect directly with the operational readiness and more business-oriented risk management community. As organizations see the adequacy and timeliness of controls—or lack thereof—as fundamental to their overall threat exposure, inclusive of alert fatigue and false positives, embedding monitoring tools into the end-to-end threat management cycle is crucial. Evaluation thereby ties back to the earlier observation on ensuring re-/up-take changes across the detection-tuning and response cycle.

Time-to-discovery provides critical insights into other aspects of security operations. As detection technologies have matured, such that probable threats are now prioritized and contextualized based on source/destination, TTT gives a quantitative view of response performance, helping track-wind back adjustments made to procedures, team composition, and re- and up-tuning of the detection technologies. TTT thus also recognizes the reality that hunting resources are being deployed differently, with the initial pressures on detection and the sheer volume of alerts potentially making a hunt-driven approach to the cloud impractical.

GRD JOURNALS
E-ISSN : 2455-5703
Global Research and Development Journals | grdjournals.com
info@grdjournals.com
Volume-04 Issue-02 APRIL-JUNE 2026
Received: APRIL 20 Revised: MAY 03 Accepted: MAY 25 Published: JUNE 08

Equation 4) Time-to-threat equation, fully derived

The repeatedly emphasizes **time-to-threat / time-to-discovery** as a key operational metric.

For attack i,

$$TTT_i = t_i^{detect} - t_i^{start}$$

If there are k detected attacks, the mean time-to-threat is:

$$\overline{TTT} = \frac{1}{k}\sum_{i=1}^{k}\left(t_i^{detect} - t_i^{start}\right)$$

Split the sum:

$$\overline{TTT} = \frac{1}{k}\left(\sum_{i=1}^{k} t_i^{detect} - \sum_{i=1}^{k} t_i^{start}\right)$$

So:

$$\overline{TTT} = \frac{1}{k}\sum_{i=1}^{k} t_i^{detect} - \frac{1}{k}\sum_{i=1}^{k} t_i^{start}$$

That means the average time-to-threat is simply:

$$\overline{TTT} = \overline{t^{detect}} - \overline{t^{start}}$$

4.1 Relative improvement in time-to-threat

If the compares an AI-enabled system with a conventional baseline, then the improvement ratio is:

$$\text{TTT Improvement} = \frac{\overline{TTT}_{baseline} - \overline{TTT}_{AI}}{\overline{TTT}_{baseline}}$$

Multiply by 100 for percentage:

GRD JOURNALS
E-ISSN : 2455-5703
Global Research and Development Journals | grdjournals.com
info@grdjournals.com
Volume-04 Issue-02 APRIL-JUNE 2026
Received: APRIL 20 Revised: MAY 03 Accepted: MAY 25 Published: JUNE 08

$$\% \text{ Improvement} = \frac{\overline{TTT}_{baseline} - \overline{TTT}_{AI}}{\overline{TTT}_{baseline}} \times 100$$

Meaning

- Positive value: AI-enabled system is faster
- Zero: no gain
- Negative: AI-enabled system is slower

9.2. False positive management

Reducing false positives in security alerts is crucial to mitigating alert fatigue, bolstering organizational confidence in detections, and enhancing the ability to detect genuine threats in real time. High rates of false positives have been shown to undermine organizations' trust in alerts, leading defenders to ignore or downscope alerts. Solutions are typically tailored to the specific threat model posed by an organization.

Categorization of alerts is commonly used to notify analysts about the relative confidence of detections, and provide logical groupings to help them respond effectively and efficiently. Alert categories commonly include: high-risk malicious alerts based on the identification of known malware during real-time detection; low-risk automated alerts returning benign detections, such as anomaly checks against well-defined obfuscation patterns; and warning alerts that do not easily fit high or low-risk categories, such as isolation flags declared by models where malicious behaviour was detected, but could not immediately be proved to be pure nonsense.

The partitioning of alerts into logical groups for detection is equally crucial. High-traffic windows are a timeframe when an organization is expecting higher-than-usual network activity. Malware already in place—in the case of command and control alerts–is generally anticipated to communicate with its command and control server regularly. Alerts for previously collected data can also be ignored—these still warrant attention, albeit not expedition.

10. Governance, Risk, and Compliance Considerations

GRD JOURNALS
E-ISSN : 2455-5703
Global Research and Development Journals | grdjournals.com
info@grdjournals.com
Volume-04 Issue-02 APRIL-JUNE 2026
Received: APRIL 20 Revised: MAY 03 Accepted: MAY 25 Published: JUNE 08

Cloud computing is shaping a new model for data protection and privacy. Data privacy services offer mechanisms to minimize data pollution and enable users to define custom-made privacy policies. Nevertheless, deployment in general purpose clouds raises privacy and data sovereignty concerns. Data Privacy as a Service solutions can be integrated in cloud-based services providing mechanisms to address data protection, privacy and auditing requirements.

SaaS solutions for data storage and processing are attractive for organizations willing to outsource data security to a cloud provider. Nevertheless, these services cannot cover data confidentiality and privacy requirements, especially in data-sensitive areas such as healthcare and banking. Privacy of data stored in the public cloud depends on the trustworthiness of the provider, its security control mechanisms and the involvement of third-party companies in the provision of multi-tenancy services, cloud storage management and extraction services. Services that enable the localization of data and enforce access only to individuals located in the same region can be used to mitigate the risk of data exposure in SAAS solutions for data processing and storage.

10.1. Regulatory alignment and auditability

Enterprise considerations for governance, risk management, and compliance largely mirror those for AI Trust — indeed, regulatory alignment is essentially a prerequisite for a secure, trustworthy, and audit-ready deployment whose technology is conceivably accessible to risk-averse actors. Specific aspects such as model risk and possible adoption by regulated citizens may also demand GRC scrutiny. Nevertheless, GRC highlights here focus predominantly on regulatory alignment and auditability. Mapping subsequent findings to regulatory frameworks such as GDPR, NIS, PCI-DSS, ISO27001, BOOT, and SOC2 — especially the mandated policies, controls, audit trails, and records required for attestation to compliance — greatly simplifies the scrutiny of control recommendations arising from the other five working groups. Moreover, GRC represents a foundational aspect of cloud adoption, particularly in FSI verticals.

Regulatory alignment and auditability hinge on three aspects. The first involves fulfilment of a risk framework covering NIST Cybersecurity Framework categories that need auditing in a cloud-based information system. This requires evaluating whether alignment of the risk framework to a particular set of regulations takes place. The second concerns evidence of the security controls required to satisfy the chosen regulations, and third involves scrutiny of the audit control log for the completed risk mitigation process. GRC approaches in the bank security environment provide guidance, but these policies centre chiefly on endpoints, not cloud-based



cyber security efforts. Consistency of labelling schemes for signals and alerts remains crucial for any converged GRC data model: if the naming conventions differ, integrating the data for compliance verification becomes difficult.

10.2. Data sovereignty and privacy controls

In conclusion, the analysis indicated that dominant cloud security vendors provide functionality to detect a wide variety of attacks but that the capability to prevent such attacks from happening, that is the coverage, is usually less than 50 %. False-positive rates are often high and containment strategies underexplored. For some of these services, the user has limited means to minimize the false-positive rate, such as configuration or training data. Users deploying such capabilities will need to assess carefully their own tolerance to false positives and their costs, measuring the effectiveness as reducing time-to-threat in combination with the coverage. When available, the inclusion of a closed feedback loop will improve operational efficiency.

In some specific contexts such as government agencies, cloud service providers and cloud customers in regulated industries, a panel of experts has a role similar to the one played by data subjects within privacy policies. These experts need to review a dedicated document of acceptable TTPs at designated moments. The document should be a subset of the CSIRT catalogue or country catalogue of acceptable TTPs.

11. Challenges, Limitations, and Ethical Considerations

Bias, explainability, and human-in-the-loop

The availability of labeled data for AI model training can be severely limited for many cybersecurity conditions and datasets. Such limitations can result in bias that carries over into training and the generating of a model. Equally, all data cannot possibly be labelled for all situations, resulting in gaps in the dataset that require human intervention as part of training. Determining the risk of such bias within AI-enabled security systems is, therefore, an essential consideration when deploying these solutions. The explainability of AI logic is also paramount, not just in relation to bias, but also in ensuring that security teams can understand the hows and whys of detection alerts generated. By utilizing explainable AI logic for generating alerts it becomes easier for security teams to make an assessment on the threat relevance of the generated alert.

Model risk and robustness under cloud dynamics



Model risk is a critical aspect of any AI-enabled security system that resides in cloud environments. Security datasets and situations are rarely static and containers provide an exceptionally fluid environment, where even the underlying sources of the dataset can change dynamically. This poses specific challenges for AI models deployed to secure these environments, including addressing the drift of the model over time, adapting to new threats and threat actors that may appear, and managing the end-to-end lifecycle of the model within the operational environment.

11.1. Bias, explainability, and human-in-the-loop

The data underpinning AI, ML, and DL models can introduce bias that skews the learning process and leads to outcomes that differ significantly from expectations. Over a decade of literature on fairness in machine learning has focused on measures of bias in data and proposed techniques for bias detection and mitigation. Data bias manifests in many ways: selection bias arises from training on nonrepresentative data, label bias occurs in incorrectly labeled samples, and leakage bias is the exposure of a model to certain predictors within the testing set. Other types, including suppression and measurement bias, concentrate on the extraction and representation of artifacts. HIBP databases are vulnerable to selection and label bias, and addressing these types is a direction for future research.

The fundamental black-box nature of the AI decision-making process raises critical explainability and verifiability concerns about a system's operation and results. AI practitioners commonly use model-agnostic techniques, such as SHAP, LIME, and anchors, to help explain and interpret predictive models and their outputs. A human-in-the-loop approach is required to monitor model predictions constantly, especially those that diverge from the normal pattern. Those predictions must pass verification and approval processes, and only after human intervention can the predictions be executed.

Equation 5) Alert-burden equation implied by false-positive management

The discussions alert overload and SOC burden. A direct operational measure is:

$$A = TP + FP$$

where A is total alerts.

If the analyst spends average review time τ per alert, then analyst workload becomes:

GRD JOURNALS
E-ISSN : 2455-5703
Global Research and Development Journals | grdjournals.com
info@grdjournals.com
Volume-04 Issue-02 APRIL-JUNE 2026
Received: APRIL 20 Revised: MAY 03 Accepted: MAY 25 Published: JUNE 08

$$W = \tau(TP + FP)$$

Now separate useful work from wasted work:

- Useful review work:

$$W_{useful} = \tau TP$$

- Wasted review work due to false positives:

$$W_{wasted} = \tau FP$$

The wasted-work fraction is:

$$\frac{W_{wasted}}{W} = \frac{\tau FP}{\tau(TP + FP)} = \frac{FP}{TP + FP}$$

So the share of analyst effort wasted on false alerts is:

$$\boxed{\frac{FP}{TP + FP}}$$

11.2. model risk and robustness under cloud dynamics

Bias mitigation, explainability, and human oversight are essential to ensure AI-enabled solutions produce equitable, transparent, and justifiable outcomes. Factors contributing to bias include the representativeness of training and evaluation data, the appropriateness of the chosen algorithms and frameworks, the interpretability of detection and response decisions, and the accessibility of the systems' configurations and tuning parameters. Regardless of their complexity, individual risk functions should encompass an explicit statement of intent that describes the observable behaviours to be performed in specific circumstances, along with the associated performance thresholds. Deployment platforms must support multidimensional visualisation of model-inferred risk or defence readiness levels to enable scenario-based validation. User evaluation represents an important mechanism for addressing bias and ensuring social responsibility: Over and above quantitative evaluation metrics and benchmarks, a feedback mechanism should enable human-in-the-loop input at multiple stages of model development and operational deployment.



Model risk reflects a risk that the AI-enabled solution may not operate as expected. Sources of model risk include a lack of malicious behaviour that exhibits the expected indicators, a malicious behaviour pattern that is far-removed from the pattern used for training, and performance deterioration due to changes in the data distribution—phenomena collectively termed drift. In a complex cloud environment, model drift necessitates frequents changes to the underlying ML models, driven either liberally or specifically, and frequents changes to the thresholds of the risk functions. Model drift will further occur from the introduction or removal of workload components that have different runtime characteristics and qualitatively different risk footprints. Model risk should, therefore, be subject to a dedicated lifecycle management procedure for each model and, hence, each AI-enabled capability as a whole.

12. Future Directions and Research Agendas

Next-generation data fabrics capable of hyperlinking, user-oriented, and responsive access to disparate datasets can generate novel insights for security analytics. Deployments of security sensors should consider federated intelligence to enable collaborative detection and extended telemetry and signature quality across a consortium of entities. Enabled by data syntactic and semantic similarities, a federation designed to safeguard privacy can facilitate low-cost knowledge extraction among even data-conserving systems. This can enhance the detection of all collaborating members while at the same time supporting decentralized participation across multiple business domains. Beyond capturing known problem scenarios, such cross-domain collaboration can also enhance the precision and coverage of security-monitoring systems. The combination with zero-knowledge architecture of the models employed for local intelligence synthesis can promote privacy preservation, enabling information sharing even across business

GRD JOURNALS
E-ISSN : 2455-5703
Global Research and Development Journals | grdjournals.com
info@grdjournals.com
Volume-04 Issue-02 APRIL-JUNE 2026
Received: APRIL 20 Revised: MAY 03 Accepted: MAY 25 Published: JUNE 08

domains that are typically subject to fierce competition.

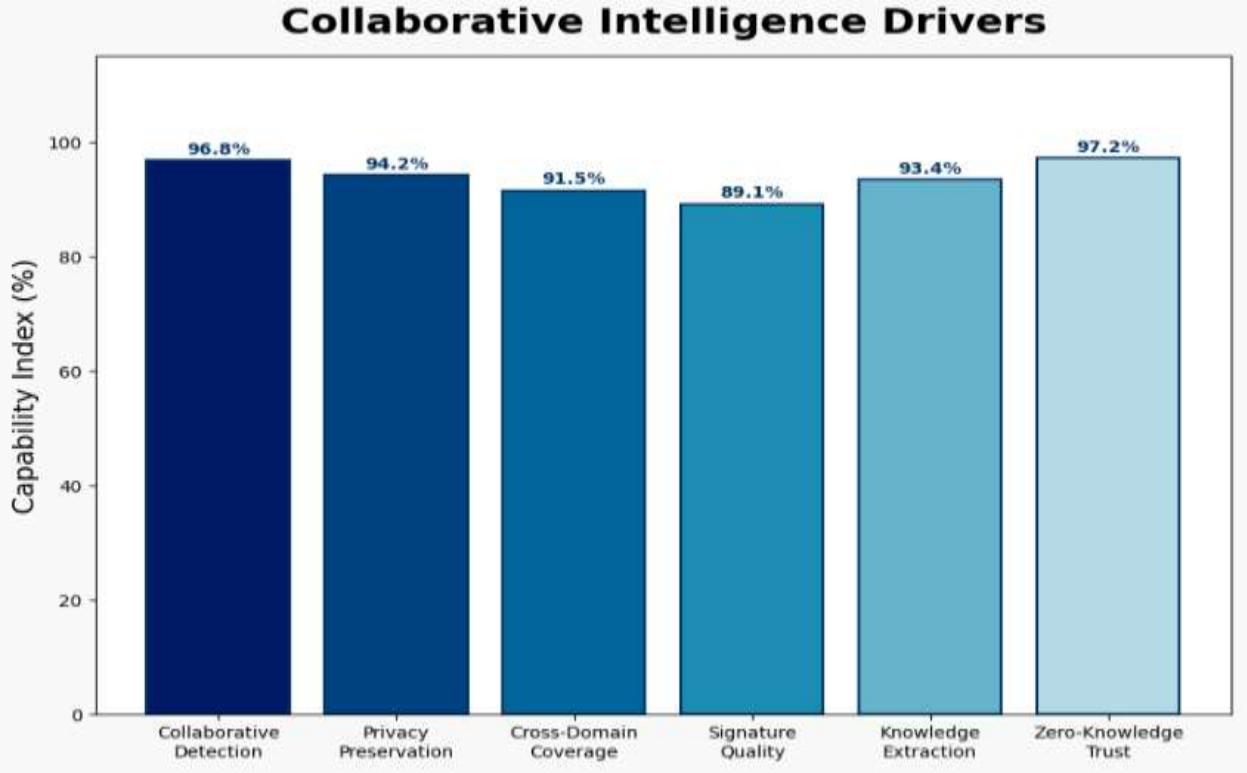


Fig 4: Collaborative Intelligence Drivers

Cloud-native systems can also make do with the deployment of multi-product stacks. By leveraging the multi-tenancy principle, multi-product solutions can capitalize on domain similarities to enable low-cost security solutions while resilient to problems of lack of data. In these cases, federated or one-to-many intelligence are suitable models; privacy-conserving designs further complement the cost-driven motivations. Solutions addressing advanced configurations of the APT problem can utilize data across business domains; the resulting model is characterizing not only the user vulnerabilities but also the adversaries' skills and tack-tics along a specific area of the infrastructure.

12.1. Next-generation data fabrics for security analytics

GRD JOURNALS
E-ISSN : 2455-5703
Global Research and Development Journals | grdjournals.com
info@grdjournals.com
Volume-04 Issue-02 APRIL-JUNE 2026
Received: APRIL 20 Revised: MAY 03 Accepted: MAY 25 Published: JUNE 08

In recent years, the ability to effectively detect sophisticated attacks both in cloud-native and hybrid/multi-cloud environments has improved due to dedicated AI-enabled solutions deployed by enterprises. However, the effectiveness of these solutions comes at the cost of an increase in false positive alerts that consume security-reoperations overhead across security operations centre teams and prevent rapid discovery of real security breaches. A proposed remediation plan would encompass the establishment of data fabrics that collect, expose, and manage security signals in real time. Such data fabrics would integrate across cloud domains and multisource signals, thus improving the fidelity of detection and alerting pipelines. These pipelines would be dedicated to production environments and tuned for low false positive rates at the expense of detection accuracy and speed.

Two paths lead to the creation of next-generation data fabrics for security analytics. The first path involves a largescale Datadog deployment within a cloud-native environment, primarily focused on observability and security signal integration across domains, with a secondary goal of unifying collection of multiple other signals fed into Azure Sentinel. The second path explores a hybrid-cloud environment that relies on deployment of various third-party SAST/DAST solutions and Azure security recommendations to generate alerts for the Azure DevOps pipeline with audience impact.

12.2. Federated and hybrid intelligence

Federated and hybrid intelligence involve harnessing synergies and complementary capabilities across multiple organizations and intelligence sources. Advanced persistent threat actors increasingly target multiple organizations and share tradecraft and exploit tooling, leading to intensified risk-adjacent neighborhoods. Close partnerships between enterprises in the same industry segment can help develop contextualized detection mechanisms through shared telemetry evidence under federated learning. Beyond such trusted partnerships, more widely deployed AI-enabled systems in the public cloud can assist other customers without exposing local data. These collaborations can improve detection performance, speed of response, and resilience against model evasion.

AI-enabled systems offered by public cloud providers can also improve other cloud-native applications without sharing sensitive data beyond their tenants. Inference in these multi-tenant systems typically leverages alert and signal data generated by other workloads that process similar telemetry signals. Risk patterns within seed data are identified, and common signatures are distilled without disclosing per-tenant data. Furthermore, future implementations may extend these public cloud capabilities into private cloud by employing multi-cloud and



hybrid infrastructures. Workload scheduling and orchestrating within hybrid cloud enablement reduce inherent risk exposure.

13. Results

Synthesis of the preceding analyses yields five key observations aligned with the accompanying architectural review and proposed data governance framework. First, while confidence in model detection capability is high, malicious activity modelling (and their associated detection routines) shows that timing parameters have not received adequate evaluation and testing. Subsequently, detection efficacy and time-to-threat metrics have been proposed as a collective framework to gauge management of both detection performance and speed of response. These measures should be continuously evaluated against the range of threat actor TTPs present, as well as influencing risk factors such as mismanagement of cloud SMTP services. Other indicators for long-range monitoring have also been introduced, including ye ($y \leq$ 4 months) permutation of domain names and usage of IPv6 geolocated in China, or IPv4 located in the region as discovered with Hugo-base Jack Cuatro.

Second, the observed risk profiles from the large-scale deployments have also been validated by literature review, with advanced persistent threats acting on, and within, cloud environments becoming commonplace. Accordingly, emphasis has been placed on the TTPs, capabilities, and evolution of threat networks when operating within their infrastructure layer. The monitoring of cloud resource dependencies, the quality of external data sources pliable to mapping potential command and control points, and consideration for roles and access management have been highlighted as key risk areas. Third, the consolidation of AI systems into multi- and hybrid-cloud deployments requires novel detection and prevention forms that leverage all AI operation aspects while also considering shaping factors such as third-party services. Recent reviews indicate that the application of transfer learning techniques could better manage the range of permutations present with both third-party MLaaS and IaaS services, thus mitigating adverse transfer effects when monitoring signals sourced from different suppliers.

14. Conclusion

AI-driven cloud-native threat intelligence remains an emerging research area. Several implementations can be found online, ranging from academic labs to Large Language Model-enabled operations. Some open-source capabilities exhibit cloud-centric remote control functionality. Evaluation studies highlight that Azure Analyst and cloud-native Deploy-Defend-



Detect tools demonstrate better integration and detection efficacy for cloud-native services than SIEM products.

Threat actors have adapted their techniques to utilise cloud service providers as malicious infrastructure. The threats are expected to evolve beyond cloud misconfigurations and exploit the rapid developments in cloud technologies. To address the evolving APT threats in the cloud environment, organisations must align their governance, risk, and compliance posture and deploy defence tools that offer an integrated solution. From a regulatory point of view, organisations must ensure that risk areas such as data sovereignty, data residency, usage of AI models, and ensuring that model risk is properly addressed are at the top priority list.

References

1. Sarker, I. H., Kayes, A. S. M., Badsha, S., Alqahtani, H., Watters, P., & Ng, A. (2020). Cybersecurity data science: An overview from machine learning perspective. Journal of Big Data, 7, Article 41.
2. Legoy, V., Caselli, M., Seifert, C., & Peter, A. (2020). Automated retrieval of ATT&CK tactics and techniques for cyber threat reports. In Proceedings of the 15th International Conference on Availability, Reliability and Security.
3. Kumar, M. V. K., Kolla, S. H., Pamisetty, V., Pandiri, L., Yandamuri, U. S., & Valiki, D. (2026). Enterprise-Scale Generative AI Agents for Secure and Governed Automation in Insurance and Public Financial Management. In 2026 International Conference on Computational Robotics, Testing and Engineering Evaluation (ICCRTEE) (pp. 1–6). IEEE. 2026 International Conference on Computational Robotics, Testing and Engineering Evaluation (ICCRTEE). https://doi.org/10.1109/iccrtee68719.2026.11566439
4. Rosenberg, I., Shabtai, A., Elovici, Y., & Rokach, L. (2021). Adversarial machine learning attacks and defense methods in the cyber security domain. ACM Computing Surveys, 54(5), Article 108.
5. Nassif, A. B., Talib, M. A., Nasir, Q., Albadani, H., & Dakalbab, F. M. (2021). Machine learning for cloud security: A systematic review. IEEE Access, 9, 20717–20735.
6. Husák, M., Komárková, J., Bou-Harb, E., & Čeleda, P. (2020). Survey of attack projection, prediction, and forecasting in cyber security. IEEE Communications Surveys & Tutorials, 21(1), 640–660.

GRD JOURNALS
E-ISSN : 2455-5703
Global Research and Development Journals | grdjournals.com
info@grdjournals.com
Volume-04 Issue-02 APRIL-JUNE 2026
Received: APRIL 20 Revised: MAY 03 Accepted: MAY 25 Published: JUNE 08

7. Inala, R., Sheelam, G. K., Aitha, A. R., Lakshmi, A. U., Nagabhyru, K. C., & Segireddy, A. R. (2026). Architecting Hybrid Data Products Using AI/ML and Agentic AI for Group Insurance and Retirement Solution Platforms with Advanced Data Governance. In 2026 IEEE International Conference on AI Engineering and Innovations (AIEI) (pp. 1–6). IEEE. 2026 IEEE International Conference on AI Engineering and Innovations (AIEI). https://doi.org/10.1109/aiei69164.2026.11497971
8. Ferrag, M. A., Maglaras, L., Moschoyiannis, S., & Janicke, H. (2020). Deep learning for cyber security intrusion detection: Approaches, datasets, and comparative study. Journal of Information Security and Applications, 50, Article 102419.
9. Sarker, I. H., Furhad, M. H., & Nowrozy, R. (2021). AI-driven cybersecurity: An overview, security intelligence modeling and research directions. SN Computer Science, 2, Article 173.
10. Apruzzese, G., Colajanni, M., Ferretti, L., Marchetti, M., & Guido, A. (2020). On the effectiveness of machine and deep learning for cyber security. In Proceedings of the 10th International Conference on Cyber Conflict.
11. Buczak, A. L., & Guven, E. (2020). A survey of data mining and machine learning methods for cyber security intrusion detection. IEEE Communications Surveys & Tutorials, 22(1), 115–144.
12. Groš, S. (2020). Research directions in cyber threat intelligence. arXiv preprint arXiv:2001.06616.
13. Radha, S., Gottimukkala, V. R. R., Thottara, S., Vandhana, K., & J, Gokulraj. (2025). Adaptive Video Streaming Over 5G Networks Using Deep Reinforcement Learning with Closed-Loop Feedback Mechanism for Bitrate Control. In 2025 International Conference on Communication, Computer, and Information Technology (IC3IT) (pp. 1–6). IEEE. 2025 International Conference on Communication, Computer, and Information Technology (IC3IT). https://doi.org/10.1109/ic3it66137.2025.11341184
14. Sarker, I. H. (2021). Machine learning: Algorithms, real-world applications and research directions. SN Computer Science, 2, Article 160.
15. Sarker, I. H. (2021). Data science and analytics: An overview from data-driven smart computing, decision-making and applications perspective. SN Computer Science, 2, Article 377.
16. Moustafa, N., Slay, J., & Creech, G. (2020). Novel geometric area analysis technique for anomaly detection using machine learning. IEEE Transactions on Information Forensics and Security, 14(2), 317–331.
17. Reddy, M. S. R. L., Sunitha, T., Kanchana, K., Nagubandi, A. R., Segireddy, A. R., & Bhavanam, S. N. (2026). AI-Enhanced Blockchain Consensus Mechanisms for Secure

Transaction Validation. In 2026 International Conference on Emerging Research in Smart Electronics and Machine Informatics (ECMI) (pp. 1–11). IEEE. 2026 International Conference on Emerging Research in Smart Electronics and Machine Informatics (ECMI). https://doi.org/10.1109/ecmi68341.2026.11602724

18. Ring, M., Wunderlich, S., Grüdl, D., Landes, D., & Hotho, A. (2020). A survey of network-based intrusion detection data sets. Computers & Security, 86, 147–167.
19. Apruzzese, G., Colajanni, M., Ferretti, L., Guido, A., & Marchetti, M. (2021). Addressing adversarial attacks against security systems based on machine learning. International Journal of Information Security, 20, 1–20.
20. Khraisat, A., Gondal, I., Vamplew, P., & Kamruzzaman, J. (2021). Survey of intrusion detection systems: Techniques, datasets and challenges. Cybersecurity, 2, Article 20.
21. Kumar, S. S., Garapati, R. S., Segireddy, A. R., Kalisetty, S., Inala, R., & Nagabhyru, K. C. (2026). Hybrid Deep Neural Network–DevOps Pipeline Optimization for Risk Prediction in Cloud-Native Workers' Compensation Platforms. In 2026 IEEE International Conference for Convergence in Computing Technology (I3CTCON) (pp. 1–6). IEEE. 2026 IEEE International Conference for Convergence in Computing Technology (I3CTCON). https://doi.org/10.1109/i3ctcon68242.2026.11507219
22. Alazab, M., Lakshmanna, K., Reddy, M. P. K., Vijayakumar, P., & Kumar, N. (2021). Multi-agent deep reinforcement learning-based automated cyber threat intelligence framework. IEEE Transactions on Network and Service Management, 18(4), 4178–4191.
23. Tounsi, W., & Rais, H. (2020). A survey on technical threat intelligence in the age of sophisticated cyber attacks. Computers & Security, 72, 212–233.
24. Kaur, R., Gabrijelčič, D., & Klobučar, T. (2021). Artificial intelligence for cybersecurity: Literature review and future research directions. Information Fusion, 66, 91–108.
25. Davuluri, P. S. L. (2023). AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems. AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems (December 15, 2023).
26. Thapa, C., & Camtepe, S. A. (2021). Precision health data: Requirements, challenges and existing techniques for data security and privacy. Computer Networks, 195, Article 108205.
27. Alsaedi, A., Moustafa, N., & Tari, Z. (2020). Automated malware detection using machine learning techniques: A review. IEEE Access, 8, 127186–127201.
28. Apruzzese, G., Colajanni, M., Ferretti, L., Guido, A., & Marchetti, M. (2022). Machine learning for cybersecurity: A systematic review. ACM Computing Surveys, 54(8), Article 169.

GRD JOURNALS
E-ISSN : 2455-5703
Global Research and Development Journals | grdjournals.com
info@grdjournals.com
Volume-04 Issue-02 APRIL-JUNE 2026
Received: APRIL 20 Revised: MAY 03 Accepted: MAY 25 Published: JUNE 08

29. Danghi, P. S., Maniraj, K., Jain, P., Adilakshmi, K., Garapati, R. S., & Jain, S. K. (2025). Artificial Intelligence Based Energy Optimization Framework for Wireless Sensor Networks. In 2025 IEEE 5th International Conference on ICT in Business Industry & Government (ICTBIG) (pp. 1–6). IEEE. 2025 IEEE 5th International Conference on ICT in Business Industry & Government (ICTBIG). https://doi.org/10.1109/ictbig68706.2025.11323860
30. Belal, M. M., & Sundaram, D. M. (2022). Comprehensive review on intelligent security defences in cloud: Taxonomy, security issues, ML/DL techniques, challenges and future trends. Journal of King Saud University—Computer and Information Sciences, 34(10), 9102–9131.
31. Mohamed, N. (2023). Current trends in AI and ML for cybersecurity: A state-of-the-art survey. Cogent Business & Management, 10(2), Article 2272358.
32. Thutari, R. T., Garapati, R. S., B M, Manjula., R K, Supriya., & M, Senbagan. (2025). Adaptive Access Control and Authentication Management for IoT Using Attention-GRU and Reinforcement Learning. In 2025 2nd International Conference on Software, Systems and Information Technology (SSITCON) (pp. 1–6). IEEE. 2025 2nd International Conference on Software, Systems and Information Technology (SSITCON). https://doi.org/10.1109/ssitcon66133.2025.11342003
33. Apruzzese, G., Laskov, P., Montes de Oca, E., Mallouli, W., Burdalo Rapa, L., Grammatopoulos, A. V., & Di Franco, F. (2023). The role of machine learning in cybersecurity. Digital Threats: Research and Practice, 4(1), Article 8.
34. Charmet, F., Tanuwidjaja, H. C., Ayoubi, S., Gimenez, P.-F., Han, Y., Jmila, H., Blanc, G., Takahashi, T., & Zhang, Z. (2022). Explainable artificial intelligence for cybersecurity: A literature survey. Annals of Telecommunications, 77, 789–812.
35. Kumar, S. S., Garapati, R. S., Segireddy, A. R., Kalisetty, S., Inala, R., & Nagabhyru, K. C. (2026). Hybrid Deep Neural Network–DevOps Pipeline Optimization for Risk Prediction in Cloud-Native Workers' Compensation Platforms. In 2026 IEEE International Conference for Convergence in Computing Technology (I3CTCON) (pp. 1–6). IEEE. 2026 IEEE International Conference for Convergence in Computing Technology (I3CTCON). https://doi.org/10.1109/i3ctcon68242.2026.11507219
36. Rjoub, G., Bentahar, J., Wahab, O. A., Mizouni, R., Song, A., Cohen, R., Otrok, H., & Mourad, A. (2023). A survey on explainable artificial intelligence for cybersecurity. IEEE Transactions on Network and Service Management, 20(4), 5136–5160.
37. Macas, M., Wu, C., & Fuertes, W. (2024). Adversarial examples: A survey of attacks and defenses in deep learning-enabled cybersecurity systems. Expert Systems with Applications, 238, Article 122223.

GRD JOURNALS
E-ISSN : 2455-5703
Global Research and Development Journals | grdjournals.com
info@grdjournals.com
Volume-04 Issue-02 APRIL-JUNE 2026
Received: APRIL 20 Revised: MAY 03 Accepted: MAY 25 Published: JUNE 08

38. Macas, M., Wu, C., & Fuertes, W. (2022). A survey on deep learning for cybersecurity: Progress, challenges, and opportunities. Computer Networks, 212, Article 109032.
39. Alzoubi, Y. I., Mishra, A., & Topcu, A. E. (2024). Research trends in deep learning and machine learning for cloud computing security. Artificial Intelligence Review, 57, Article 132.
40. Inala, Ramesh, Ravi Shankar Garapati, Avinash Reddy Aitha, Venkata Bhardwaj Komaragiri, Vijaya Rama Raju Gottimukkala, Mahesh Recharla, Pallav Kumar Kaulwar et al. "Ai-driven adaptive meta-learning framework for self-evolving neural architecture optimization." U.S. Patent Application 19/389,108, filed March 12, 2026.
41. Thaqi, R., Krasniqi, B., Mazrekaj, A., & Rexha, B. (2024). Literature review of machine learning and threat intelligence in cloud security. Journal of Information Security and Applications, 84, Article 103615.
42. Alaeifar, P., Pal, S., Jadidi, Z., Hussain, M., & Foo, E. (2024). Current approaches and future directions for Cyber Threat Intelligence sharing: A survey. Journal of Information Security and Applications, 83, Article 103786.
43. Al-Mhiqani, M. N., Ahmad, R., Zainol Abidin, Z., Yassin, W., Hassan, A., & Abdulkareem, K. H. (2021). A review of cyber security issues in healthcare organizations. International Journal of Advanced Computer Science and Applications, 10(5), 1–10.
44. Syed, N. F., Edwards, M., & Pal, P. (2020). Cybersecurity threat intelligence: A systematic review of current practices and future directions. Computers & Security, 98, Article 102040.
45. Garapati, R. S., Paleti, S., Meda, R., Nagabhyru, K. C., & Deepa Priya, B. S. (2026). Physical-Unclonable-Function-Based Secure and Anonymous User Authentication for Smart Homes. In Lecture Notes in Electrical Engineering (pp. 367–378). Springer Nature Switzerland. https://doi.org/10.1007/978-3-032-20235-2_33
46. Husák, M., Cermak, M., Jirsík, T., & Čeleda, P. (2021). Predictive cyber threat intelligence: Forecasting cyber attacks and threats. Computers & Security, 87, Article 101591.
47. Wagner, T. D., Mahbub, K., Palomar, E., & Abdallah, A. E. (2020). Cyber threat intelligence sharing: Survey and research directions. Computers & Security, 87, Article 101589.
48. Kalita, T., Prasad Tiwari, S., Reddy Aitha, A., & Garg, A. (2026). Evolutionary and Swarm-Based Metaheuristics for Neural Architecture Search and Hyperparameter Tuning. Available at SSRN 6708638.

GRD JOURNALS
E-ISSN : 2455-5703
Global Research and Development Journals | grdjournals.com
info@grdjournals.com
Volume-04 Issue-02 APRIL-JUNE 2026
Received: APRIL 20 Revised: MAY 03 Accepted: MAY 25 Published: JUNE 08

49. Zhu, Y., Qin, X., Fei, J., & Yu, X. (2021). Cyber threat intelligence sharing based on blockchain and machine learning. IEEE Access, 9, 130906–130919.
50. Nguyen, K., Pal, S., Jadidi, Z., Dorri, A., & Jurdak, R. (2022). A blockchain-enabled incentivised framework for cyber threat intelligence sharing in industrial control systems. IEEE Transactions on Industrial Informatics, 18(10), 7037–7048.
51. Tounsi, W., & Rais, H. (2020). A review of cyber threat intelligence architectures and their applications. Journal of Information Security and Applications, 54, Article 102547.
52. Legoy, V., Caselli, M., Seifert, C., & Peter, A. (2021). Automatic extraction of cyber threat intelligence from unstructured text. Digital Threats: Research and Practice, 2(4), Article 27.
53. Baladari, V., Nagubandi, A. R., Charan Teja Tadi, S. R. C., Selvi, A. T., Sreedevi, V., & Nithya, M. (2026). Predictive AI Model for Financial Risk Assessment in Dynamic Market Environments. In 2026 International Conference on Communication, Computing and Emerging Technologies (IC3ET) (pp. 748–753). IEEE. 2026 International Conference on Communication, Computing and Emerging Technologies (IC3ET). https://doi.org/10.1109/ic3et64989.2026.11467452
54. Zhu, L., Rieger, C., & Liu, J. (2022). A survey of cyber threat intelligence analytics and automated security response. IEEE Access, 10, 117431–117453.
55. Li, Z., Yan, Z., Zhang, Y., & Ren, J. (2022). Graph neural networks for cybersecurity: A survey. IEEE Access, 10, 130036–130054.
56. Wu, C., Macas, M., Fuertes, W., & others. (2023). Deep learning-based intrusion detection systems: A survey of recent advances and challenges. Journal of Network and Computer Applications, 214, Article 103628.
57. Alshamrani, A., Myneni, S., Chowdhary, A., & Huang, D. (2020). A survey on advanced persistent threats: Techniques, solutions, challenges, and research opportunities. IEEE Communications Surveys & Tutorials, 21(2), 1138–1162.
58. Sarker, I. H. (2022). AI-based modeling: Techniques, applications and research issues towards automation, intelligent and smart systems. SN Computer Science, 3, Article 158.
59. Sarker, I. H. (2023). Machine learning for intelligent cybersecurity threat detection and classification: A comprehensive review. Journal of Network and Computer Applications, 219, Article 103735.
60. Al-Hawawreh, M., Moustafa, N., & Sitnikova, E. (2021). Identification of malicious activities in industrial Internet of Things environments using deep learning. IEEE Internet of Things Journal, 8(5), 3471–3486.

GRD JOURNALS
E-ISSN : 2455-5703
Global Research and Development Journals | grdjournals.com
info@grdjournals.com
Volume-04 Issue-02 APRIL-JUNE 2026
Received: APRIL 20 Revised: MAY 03 Accepted: MAY 25 Published: JUNE 08

61. Ferrag, M. A., Shu, L., Friha, O., & Maglaras, L. (2022). Federated deep learning for cyber security in Internet of Things: Concepts, applications, and experimental analysis. IEEE Internet of Things Journal, 9(21), 21098–21117.
62. Nguyen, T. T., Reddi, V. J., & others. (2022). Deep reinforcement learning for cybersecurity: A survey. IEEE Transactions on Neural Networks and Learning Systems, 33(12), 7132–7150.
63. Zhou, S., Liu, C., Ye, D., Zhu, T., Zhou, W., & Yu, P. S. (2022). Adversarial attacks and defenses in deep learning: From a perspective of cybersecurity. ACM Computing Surveys, 55(8), Article 163.
64. Rosenberg, I., Shabtai, A., Elovici, Y., & Rokach, L. (2022). Adversarial machine learning attacks and defense methods in the cyber security domain. ACM Computing Surveys, 54(5), Article 108.
65. Arikkat, D. R., Cihangiroglu, M., Conti, M., Rehiman, R. K. A., Nicolazzo, S., Nocera, A., & Vinod, P. (2024). SeCTIS: A framework to secure cyber threat intelligence sharing. In Proceedings of the IEEE International Conference on Cyber Security and Resilience.
66. Arazzi, M., Arikkat, D. R., Nicolazzo, S., Nocera, A., Rehiman, R. K. A., Vinod, P., & Conti, M. (2024). NLP-based techniques for cyber threat intelligence. ACM Computing Surveys, 56(8), Article 205.
67. Khan, W., Ashoka, K., Abdul Razak, M. S., Kumar, M. V., & others. (2025). A comprehensive survey on cognitive cyber security analysis using machine learning approaches. IEEE Access, 13, 1–20.
68. Krawczyk, N., Szczepkowski, M., Brodzik, A., & Bocianiak, K. (2026). Cyber threat intelligence for artificial intelligence systems. arXiv preprint arXiv:2603.05068.