ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

AWS Security Architecture and Machine Learning for APT Detection in Cloud Environments

Adeolu Opeyemi Ojo[1], Samuel Babafemi Olabisi[2]

[1] University of Greater Manchester, Deane Road, Bolton, BL3 5AB, UK

[2]Dept. of Computer Science, Sikiru Adetona College of Education, Science and Technology, Omu-Ajose, Ogun State, Nigeria

DOI: https://doi.org/10.51583/IJLTEMAS.2026.150700111

Received: 02 August 2026; Accepted: 12 August 2026; Published: 20 August 2026

ABSTRACT

Cloud environments, and Amazon Web Services (AWS) in particular, host high-value data assets and mission-critical workloads that make them attractive targets for Advanced Persistent Threat (APT) actors. Because forensic investigation techniques are applied only after a breach has already been discovered, the volume and velocity of cloud-generated telemetry make proactive, automated detection capabilities essential. This paper reviews machine learning-driven anomaly detection paradigms — supervised, unsupervised, semi-supervised, and deep learning — and examines their suitability for APT detection in AWS environments. It also reviews the AWS shared-responsibility security architecture, including Identity and Access Management (IAM), encryption services, and logging and monitoring services such as AWS Cloud Trail, AWS Config, Amazon Guard Duty, Amazon Detective, and Amazon Inspector, and considers the NIST Cyber security Framework (CSF) as a governance overlay that connects these technical capabilities to organizational risk management. Drawing on this review of the peer-reviewed and primary-source literature, the paper argues that no single detection paradigm is likely sufficient on its own, and that unsupervised and semi-supervised machine learning, combined with AWS-native security services and governed by the NIST CSF, offer a more resilient conceptual basis for cloud APT defense than any single method in isolation. On this basis, the paper proposes an Integrated Cloud APT Detection and Defense Model (ICADDM) as a conceptual architecture for researchers and practitioners, maps AWS security services against the MITRE ATT&CK Cloud Matrix, and identifies the empirical validation of the model against real cloud telemetry as the principal direction for future work.

Keywords: Advanced Persistent Threats, Machine Learning, Anomaly Detection, AWS Security Architecture, NIST Cyber security Framework, MITRE ATT&CK.

INTRODUCTION

The migration of enterprise workloads — including government, healthcare, and financial-sector systems — to cloud platforms such as Amazon Web Services (AWS) has changed the cyber security threat landscape, making cloud environments a prime target for highly skilled Advanced Persistent Threat (APT) actors who prioritize long-term infiltration over opportunistic attacks [1].

Compared with ordinary cybercrime, APTs constitute a qualitatively different threat category. These actors, often associated with nation-states or well-resourced criminal groups, use dynamic tactics, techniques, and procedures (TTPs) to infiltrate target environments, establish enduring footholds, and steal sensitive information, frequently remaining undetected for months or years [2]. In cloud computing environments, APT actors take advantage of cloud-native mechanisms — including manipulation of Identity and Access Management (IAM), theft of virtual-

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

machine credentials, sharing of storage snapshots, and tampering with logging services — creating detection challenges that traditional, on-premises security tools are not designed to handle [3].

Most recent publications on cloud APT detection concentrate on forensic investigative techniques applied after a breach has already been confirmed. Proactive, automated detection is nevertheless necessary given the volume, velocity, and variety of cloud-generated telemetry, which includes identity events, network-flow records, API call logs, and configuration changes. Machine learning (ML) offers a promising approach: by learning baseline behaviors from historical data, ML models can identify deviations suggestive of APT activity in near real time, without relying solely on known attack signatures [4].

Despite this theoretical promise, the practical application of ML-driven anomaly detection in cloud environments remains underexplored. Most ML detection studies are evaluated on generic datasets rather than cloud-specific telemetry, the integration of ML-based detection with AWS-native security services is understudied, and little empirical research examines how technical detection capabilities align with governance frameworks such as the NIST Cyber security Framework (CSF) [3], [5]. Many organizations also lack a clear, organized understanding of which AWS security services address which threat vectors, even though the AWS shared-responsibility model places substantial security obligations on cloud tenants.

To address these gaps, this paper: (i) critically reviews and compares machine learning anomaly detection paradigms — supervised, unsupervised, semi-supervised, and deep learning — with respect to their strengths, limitations, and suitability for cloud APT detection; (ii) analyzes the AWS security architecture, including IAM, encryption, and AWS logging and monitoring services; (iii) evaluates the NIST CSF as a basis for aligning ML-based detection and AWS security services with organizational risk management; (iv) maps representative AWS security services against the MITRE ATT&CK Cloud Matrix; and (v) proposes, as a conceptual synthesis rather than an empirically validated finding, an integrated detection model that combines these components into a coherent cloud APT defense architecture.

Review Approach

This paper is a narrative and conceptual review rather than a systematic review: it synthesizes peer-reviewed literature on APT detection and machine learning-based anomaly detection, primary AWS technical documentation, and current government and industry threat-intelligence reporting (including CISA advisories and Mandiant's M-Trends series), in order to build an integrated conceptual model rather than to quantify the effect of a specific intervention.

Sources were identified through targeted searches of Google Scholar, ACM Digital Library, and IEEE Xplore using combinations of the terms "APT," "cloud security," "AWS," "anomaly detection," "machine learning," and "NIST Cyber security Framework," supplemented by AWS official documentation and current (2024–2026) threat-intelligence reports for material published after the academic literature. Because no formal screening protocol, inclusion criteria, or database-level search log were maintained, the review does not claim the completeness or reproducibility of a PRISMA-style systematic review, and the resulting synthesis — including the proposed ICADDM — should be read as a structured expert synthesis and conceptual proposal, not as an empirical finding.

LITERATURE REVIEW

This section reviews Advanced Persistent Threats and their cloud-specific attack vectors, documented cloud identity-abuse incidents, machine learning paradigms for anomaly detection, AWS security architecture, and the NIST Cyber security Framework as a governance overlay.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

Advanced Persistent Threats: Features and Cloud-Specific Attack Vectors

An Advanced Persistent Threat (APT) is a persistent and complex cyber-attack, historically associated with espionage campaigns against government, defense, and critical-infrastructure targets. The term gained wide currency in the security community during the 2000s as a way of describing long-running, well-resourced intrusion campaigns that differ qualitatively from opportunistic cybercrime [2]. The National Institute of Standards and Technology defines an advanced persistent threat as an adversary with sophisticated levels of expertise and significant resources that uses multiple attack vectors — including cyber, physical, and deception — to establish and extend a foothold within an organization's information technology infrastructure, typically to exfiltrate information or undermine critical operations over an extended period, while adapting to a defender's efforts to resist it [6].

The operational structure of APT campaigns is typically multi-stage. Ghafir and Prenosil [7] divide this structure into six phases: intelligence gathering, initial point of entry, command-and-control (C2) communication, lateral movement, data discovery, and data exfiltration. This lifecycle distinguishes APTs from commodity cybercrime, which tends to focus on short-term financial gain; APT actors instead pursue long-term strategic goals — financial manipulation, disruption of infrastructure, theft of intellectual property, or espionage — and are prepared to commit significant resources over long periods to achieve them [3].

APT actors exploit cloud-native features in ways that conventional, on-premises security controls are not designed to detect. Adversaries create unauthorized administrative users, generate persistent access keys, and backdoor IAM roles to sustain stealthy persistence after an initial breach. IAM, which governs access to all AWS resources, is therefore both the principal attack surface and the principal control mechanism from an APT-defense standpoint. Data loss prevention controls can be circumvented through the snapshot-sharing mechanisms of Relational Database Service (RDS) instances and Elastic Block Store (EBS) volumes, which allow data exfiltration without conventional network communication; remote access channels can be opened by manipulating security groups; and the ability to conduct forensic investigation is compromised when Cloud Trail logging is disabled or deleted [8]. These cloud-native TTPs highlight the need for detection tooling designed specifically for cloud telemetry.

Documented Cloud Identity-Abuse and APT-Relevant Security Incidents

Table 1 summarizes documented cloud security incidents and adversary techniques that illustrate the IAM-manipulation, credential-theft, and snapshot-exfiltration patterns described above, in place of the industrial-control-system incidents used in earlier drafts of this review, which did not reflect the cloud-specific focus of the paper.

Incident / Technique	Period	Sector(s) Affected	Mechanism
Capital One data breach	2019	Financial services	A server-side request forgery (SSRF) vulnerability in a misconfigured web application firewall was used to query the AWS instance metadata service and obtain temporary credentials for an over-permissioned IAM role, which was then used to access and exfiltrate data from more than 700 S3 buckets, affecting roughly 106 million records [9].
Scattered Spider / UNC3944 campaigns	2022–present	Telecommunications, gaming and hospitality, retail, insurance, technology	Social engineering (phishing, push-bombing, SIM-swapping) is used to bypass multi-factor authentication and obtain credentials, followed by abuse of federated identity providers and cloud IAM —

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

			including AWS, Azure, and Okta — to establish persistent access and move laterally [10].
Cloud storage snapshot exfiltration (T1537)	Ongoing, cross-sector	Any AWS tenant	Adversaries with sufficient IAM privileges create a snapshot of an EBS volume or database instance and modify its sharing attributes to expose it to an attacker-controlled account, exfiltrating data without generating conventional network-based exfiltration signals [11].
Aggregate cloud intrusion trends	2024	Cross-sector	Incident-response data show phishing and stolen credentials as the leading initial-access vectors for cloud compromises, with adversaries increasingly targeting federated identity providers and single sign-on portals as centralized points of access [12].

Table 1. Documented Cloud Identity-Abuse and APT-Relevant Security Incidents

Anomaly Detection Using Machine Learning

Machine learning allows a system to identify unusual behaviors and attack patterns without relying on explicit, hand-coded detection rules — a significant paradigm shift in cyber security. Machine learning encompasses algorithms that enable computers to identify patterns, predict future events, and improve performance based on experience [4]. ML approaches offer particular benefits for APT detection: APTs often use novel TTPs that signature-based detection tools cannot recognize, whereas anomaly-based ML models can identify deviations from established behavioral baselines even where a given attack pattern has not previously been observed.

The core mechanism by which machine learning supports APT detection is anomaly detection: identifying data patterns that differ from established behavioral baselines by analyzing executable behavior, access patterns, and network-flow characteristics [3], [13]. Table 2 summarizes the statistical and machine learning-based techniques used for anomaly detection in cyber security, including the deep learning and hybrid approaches discussed further in Section 3.4.

Method	**Description**	**Application to APT Detection**
Statistical models	Identify data points using distributional assumptions (e.g., Gaussian), so that deviation beyond an expected range is flagged as anomalous.	Effective for detecting spikes in login frequency, data-transfer volume, and API call volume.
Clustering and classification	Group's data points by similarity; points that do not fit an established cluster are treated as anomalies.	Applicable to classifying IAM user-behavior patterns; unusual access clustering may indicate compromised credentials or backdoor roles.
Auto encoder-based deep anomaly detection	Multi-stage neural network architectures learn a compressed representation of normal telemetry; events with high reconstruction error are flagged as anomalous.	Effective at detecting novel APT TTPs in high-volume, high-dimensional cloud telemetry such as Cloud Trail and VPC Flow Logs [15].
Hybrid signature-and-anomaly detection	Combines rule- or signature-based detection of known indicators with anomaly-based detection of statistical deviations within a single pipeline.	Offsets the weaknesses of either approach used alone, offering broader detection coverage against both known or previously unseen cloud APT techniques [16].

Paradigm	**Mechanism**	**Strengths**	**Limitations**	**Cloud APT Suitability**

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

Supervised learning	Requires a labeled dataset to train a model; by mapping input features to output labels, the algorithm learns to differentiate anomalous from normal data.	High accuracy for recognized attack patterns; can be trained to identify specific APT technique signatures [14].	Cannot identify new or zero-day APT tactics absent from the training set; requires labeled training data.	Limited for novel APT TTPs; appropriate for identifying known IAM-misuse patterns and previously recorded exfiltration signatures.
Unsupervised learning	Learns to recognize abnormalities from data patterns without labeled training material; anomalies are data points that do not fit into known clusters.	Operates on unlabeled data; can identify new or previously undiscovered attack types without pre-existing threat intelligence.	Computationally demanding at cloud scale; prone to false positives from insufficient baseline modeling.	Strong suitability for cloud environments with evolving attack patterns; especially useful for identifying previously unseen APT TTPs in Cloud Trail logs.
Semi-supervised learning	Trains anomaly detection models by combining labeled and unlabeled data; unlabeled data improves model accuracy while labeled data directs detection of known abnormalities.	Can identify both known and unknown anomalies; better suited than purely supervised methods to variable data distributions.	May inherit the limitations of both supervised and unsupervised methods, depending on the quality and distribution of available labels.	Well suited to cloud environments with substantial unlabeled operational telemetry and incident-response records providing partial threat labeling.
Deep learning	Multi-layered neural networks automatically learn intricate representations of input data, supporting anomaly detection in high-dimensional telemetry.	Handles high-dimensional and unstructured cloud telemetry effectively; automatically identifies complex input patterns [14].	Requires large training datasets and significant processing power; interpretability of detection rationale can be challenging [14].	Well suited to large-scale AWS environments producing high-dimensional Cloud Trail, VPC Flow Log, and Guard Duty telemetry, given sufficient training data.

Table 2. Methods of Anomaly Detection for Cyber security

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

Four principal ML paradigms are applicable to APT detection in cloud environments: supervised learning, unsupervised learning, semi-supervised learning, and deep learning. Each paradigm presents distinct operational characteristics, strengths, and limitations that determine its suitability for specific cloud-security contexts. Table 3 provides a comparative analysis of these paradigms with respect to cloud APT detection requirements, drawing on the classification scheme of Saabith et al. [14].

Table 3. Machine Learning Paradigms for APT Detection in the Cloud

Among the four paradigms, unsupervised and semi-supervised learning appear most applicable to cloud APT detection scenarios, because APT actors often use novel TTPs that elude supervised models trained on previous attack fingerprints [15]. Neuschmied et al. [15] used multi-stage auto encoder models — an unsupervised deep learning approach — to identify APT attacks in network datasets; their multi-stage architecture improved detection of previously unseen APT attack types relative to single-stage baselines. Similarly, Abhinav et al. [16] proposed a hybrid intrusion detection system that combines anomaly-based and signature-based approaches using machine learning, arguing that this combination provides more complete detection coverage by offsetting the limitations of each technique applied alone.

The volume and variety of available telemetry present an operational challenge specific to cloud environments: a single AWS account may generate hundreds of thousands of Cloud Trail events per day across dozens of services. To address the resulting shortage of labeled data, Myneni et al. [17] constructed a semi-synthetic APT dataset combining artificially generated and real-world telemetry to provide a realistic basis for training and evaluating ML models. Their findings point to a crucial gap: the practicality of supervised and semi-supervised techniques remains constrained by the scarcity of high-quality, labeled, cloud-specific APT datasets.

AWS Shared-Responsibility Model and Security Architecture

AWS operates a shared-responsibility model in which customers are responsible for data, identity and access controls, application-level security, and network configuration, while AWS is responsible for the security of the underlying hardware, software, networking, and facilities that constitute the cloud itself. Understanding this division of labor is a prerequisite for assessing the strengths and limitations of AWS-native security mechanisms in APT defense [18]. To meet the security needs of security-conscious organizations, AWS provides an extensive set of security capabilities spanning storage, database, and compute security. Three categories of AWS security capability are particularly important for APT defense: identity and access management, encryption, and logging and monitoring [19].

Identity and Access Management (IAM)

AWS Identity and Access Management (IAM) is the foundational mechanism for regulating access to all AWS resources and services, governing user identities and access permissions through policies, roles, and grants. From an APT-defense standpoint, IAM is simultaneously the primary attack surface and the primary control mechanism: defenders rely on IAM policy enforcement to limit the blast radius of a successful compromise, while adversaries frequently target IAM to establish persistence, escalate privileges, and maintain covert access [20].

Encryption Services

AWS provides encryption capabilities through services such as AWS Key Management Service (KMS), which manages cryptographic keys used throughout the AWS ecosystem, including for storage systems such as Elastic Block Store (EBS) and Simple Storage Service (S3). Encryption serves two functions in APT defense: it safeguards data confidentiality in the event of unauthorized access or attempted exfiltration, and it raises the cost for adversaries seeking to monetize stolen data. However, encryption also introduces forensic complications, because an APT actor who has obtained legitimate credentials can decrypt and exfiltrate data through normal channels, effectively bypassing the protection that encryption is intended to provide [20].

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

AWS Logging and Monitoring Services

AWS provides an extensive range of logging and monitoring capabilities that support governance, compliance, and security investigation. These services provide the telemetric foundation for both forensic analysis of APT activity and ML-based anomaly detection [21]. Table 4 summarizes the AWS logging and monitoring services most relevant to APT detection, together with their primary detection function and the coverage gaps that any defense architecture built on these services must account for.

AWS Service	Primary Function	APT Detection Relevance	Coverage Gap
AWS Cloud Trail	Records API-level activity across AWS accounts, including management, data, and Cloud Trail Insights events, with routing to S3 and Cloud Watch Logs [21].	The primary forensic data source for APT investigations; can be used to detect IAM manipulation, snapshot exfiltration, and defense-evasion actions such as Cloud Trail suspension or deletion.	Immutable storage configuration and log-file validation are necessary to prevent tampering by an adversary with sufficient privileges.
AWS Config	Tracks changes to AWS resource configurations and automatically compares recorded configurations against intended states [20].	Detects unauthorized configuration changes, such as security-group or IAM policy modifications, that may indicate APT lateral movement or persistence.	Provides configuration-level visibility only; it does not record data-plane events or application-layer behavior.
Amazon Guard Duty	A threat-detection service that uses machine learning and threat-intelligence feeds to analyze VPC Flow Logs, DNS logs, and Cloud Trail events for malicious or unauthorized activity.	Provides automated, ongoing threat detection, including indicators of command-and-control traffic, crypto currency mining, unusual API call patterns, and credential compromise.	Inadequate tuning increases the risk of alert fatigue; ML models may not detect new APT TTPs that are absent from current threat-intelligence feeds.
Amazon Detective	Automatically gathers log data and provides a graphical interface for security-event analysis, establishing contextual links between events and resources.	Speeds up APT investigation by placing security findings in context and illustrating the scope of a compromise.	An analytical tool rather than a real-time detection capability; its output is only as good as the upstream detections it draws on from Guard Duty and other sources.
Amazon Inspector	Continuously scans EC2 instances, container images, and Lambda functions for software vulnerabilities and unintended network exposure, automating security assessments.	Enables proactive vulnerability remediation before exploitation and identifies exploitable flaws that APT actors could use for initial access or lateral movement.	Vulnerability-focused rather than behavior-based; it does not detect post-exploitation actions by an APT actor already present in the environment.

Table 4. AWS Logging and Monitoring Services: APT Detection Relevance

Newer AWS Security Data Services

The services in Table 4 remain the core of AWS-native detection, but recent AWS releases materially change how their output can be aggregated and queried, and any current discussion of AWS-native detection should account for them. AWS Cloud Trail Lake is a managed audit and security data lake that aggregates, immutably stores, and allows SQL-based querying of Cloud Trail activity logs across multiple accounts and regions, with a default seven-year retention period [23]. Amazon Security Lake extends this further by automatically centralizing security telemetry — Cloud Trail management events, VPC Flow Logs, Route 53 resolver logs, and findings from AWS Security Hub and other integrated services — into a single, customer-owned data lake normalized to the

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

Open Cyber security Schema Framework (OCSF), which can be queried directly with Amazon Athena or ingested by third-party detection tooling [22]. For the ML-based detection paradigms discussed in Section 3.4, these services reduce the data-engineering burden of assembling a unified, cross-account, cross-service training and inference dataset, though they do not remove the underlying cost and latency considerations: querying years of immutably stored, multi-account telemetry at scale carries non-trivial compute and storage cost, and near-real-time detection use cases must weigh the latency of data-lake ingestion pipelines against the sub-minute response times that Guard Duty's native ML detections already provide. Any organization evaluating the telemetry volumes discussed in Section 3.4 — potentially millions of events per day in a large multi-account estate — should treat this cost-latency trade-off as a first-order design constraint rather than an implementation detail.

The NIST Cyber security Framework as a Governance Overlay

The NIST Cyber security Framework (CSF) offers a comprehensive, flexible approach to establishing and improving an organization's cyber security posture, tailorable to specific enterprises, risk profiles, and cyber security objectives [24]. The CSF defines five core functions — Identify, Protect, Detect, Respond, and Recover — each of which plays a distinct role in cyber security risk management. Unlike more technically oriented frameworks such as MITRE ATT&CK or the Cyber Kill Chain, the CSF offers a systematic procedure for equating cyber security investment with organizational risk tolerance [24]. Table 5 maps cloud APT defense capabilities discussed in this paper to the five CSF functions.

NIST CSF Function	Description	Cloud APT Defense Application	Supporting AWS Service
Identify	Recognizes hazards to systems, assets, data, and capabilities relevant to cyber security.	Maintaining a current inventory of AWS assets, IAM identities, and data classifications; assessing cloud workload risk relative to APT threat characteristics.	AWS Config (asset inventory); IAM Access Analyzer (access-risk assessment).
Protect	Implements safeguards to prevent cyber-attacks [24].	Establishing network segmentation via security groups and VPC design; enforcing multi-factor authentication; encrypting data in transit and at rest; enforcing least-privilege IAM policies.	AWS IAM, AWS KMS, AWS Security Groups, VPC configurations.
Detect	Identifies anomalies, intrusions, or security events quickly using detection and monitoring tools.	Applying detection rules for identified cloud-specific attack patterns; using ML-based anomaly detection to identify APT TTPs in near real time.	Cloud Trail with Athena SQL detection scripts; unsupervised/semi-supervised ML models for behavioral anomaly detection; Amazon Guard Duty.
Respond	Plans, organizes, and coordinates incident response to mitigate the consequences of cyber security incidents.	Executing pre-defined incident-response playbooks for cloud APT scenarios; revoking compromised access keys; isolating compromised IAM identities; preserving Cloud Trail evidence.	Cloud Trail (evidence preservation); AWS Systems Manager (automated remediation); Amazon Detective (investigation support).

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

Recover	Provides backup and recovery mechanisms and restores system capabilities after a cyber-security event [24].	Restoring IAM access controls; recovering compromised AWS resources from clean snapshots; assessing and updating detection rules in light of incident analysis.	IAM role restoration, AWS Backup, and post-event analysis of Cloud Trail logs.

Table 5. Cloud APT Defense Capabilities Mapped to NIST CSF Functions

The NIST CSF is directly relevant to cloud APT protection because, although ML-based detection tools and AWS-native security services provide the technical capabilities described above, their effective implementation within an organizational security program requires the kind of governance structure the NIST CSF is specifically designed to provide [24].

Research Gaps

This section sets out the research gaps that motivate the conceptual model proposed in Section 6. The literature reviewed above points to several gaps in current cloud APT detection practice and research. First, empirical research evaluating ML-based anomaly detection methods against cloud-native telemetry — specifically AWS Cloud Trail, Guard Duty, and VPC Flow Log data — remains conspicuously scarce; most ML-based APT detection studies are evaluated on general network-intrusion datasets that do not reflect the API-centric, event-driven character of cloud telemetry [3]. Second, although individual AWS security services have been examined separately, their combined contribution to a layered APT defense architecture — including their joint performance against the tactics and techniques in the MITRE ATT&CK Cloud Matrix — remains understudied. Third, the intersection of ML-based detection, AWS-native security services, and NIST CSF governance has not been investigated empirically, leaving practitioners without an evidence-based approach for deploying these capabilities holistically [3]. Fourth, the practical deployment of supervised and semi-supervised ML techniques remains constrained by the scarcity of high-quality, labeled cloud APT datasets. Myneni et al. [17] partially addressed this gap by constructing semi-synthetic datasets, although how representative these datasets are of real-world cloud APT campaigns remains unclear. Fifth, insufficient empirical research has examined the scalability of ML-based detection in large, multi-account AWS environments, where telemetry volumes may reach millions of events per day [16]; most existing studies operate at small-to-medium dataset scale. Finally, adversarial machine learning — the intentional manipulation of ML-based detection systems by capable APT actors — remains an under-examined threat dimension in the cloud security literature.

Mapping Aws Security Services To The Mitre Att&Ck Cloud Matrix

One of the gaps identified above is the absence of an explicit mapping between AWS security services and the tactics and techniques documented in the MITRE ATT&CK Cloud (IaaS) Matrix, which organizes adversary behavior against cloud infrastructure into eleven tactics spanning the intrusion lifecycle from initial access to impact [11]. Table 6 provides an illustrative, non-exhaustive mapping between representative ATT&CK Cloud tactics and techniques, the AWS services most relevant to detecting or preventing them, and the corresponding layer of the ICADDM proposed in Section 6. This mapping is intended as a starting point for the systematic, empirically validated coverage assessment identified as a research gap in Section 4, not as a substitute for it.

ATT&CK Tactic	**Representative Technique**	**Primary AWS Service(s)**	**ICADDM Layer**
Initial Access	Exploit Public-Facing Application; Valid Accounts: Cloud Accounts (T1190, T1078.004)	Amazon Guard Duty; AWS WAF	Layer 2 – Telemetry Collection

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

Persistence	Create Account; Account Manipulation (T1136, T1098)	AWS IAM Access Analyzer; AWS Config	Layer 1 – Asset and Identity Governance
Privilege Escalation	Abuse of overly permissive IAM roles or policies	IAM Access Analyzer	Layer 1 – Asset and Identity Governance
Defense Evasion	Impair Defenses: disabling or deleting Cloud Trail logs (T1562)	Cloud Trail log-file validation; S3 Object Lock	Layer 2 – Telemetry Collection
Credential Access	Unsecured Credentials; Steal Application Access Token (T1552, T1528)	Amazon Guard Duty; AWS Secrets Manager	Layer 2/3 – Telemetry Collection / ML Detection
Discovery	Cloud Infrastructure Discovery; Cloud Service Discovery (T1580, T1526)	Amazon Guard Duty; AWS Config	Layer 2 – Telemetry Collection
Collection	Data from Cloud Storage Object (T1530)	Guard Duty S3 Protection; S3 access logging	Layer 2/3 – Telemetry Collection / ML Detection
Exfiltration	Transfer Data to Cloud Account, e.g., EBS/RDS snapshot sharing (T1537)	AWS Config rules for public snapshot sharing; Guard Duty	Layer 3 – ML-Driven Anomaly Detection
Impact	Resource hijacking or destructive action against cloud infrastructure	AWS Backup; AWS Systems Manager	Layer 4 – Incident Response and Recovery

Table 6. Illustrative Mapping of AWS Security Services to the MITRE ATT&CK Cloud (IaaS) Matrix

Proposed Integrated Cloud Apt Detection and Defense Model (Icaddm)

Building on the synthesis of ML detection paradigms, AWS security architecture, and NIST CSF governance concepts set out above, this section proposes an Integrated Cloud APT Detection and Defense Model (ICADDM) with four layers. The ICADDM is offered as a conceptual synthesis intended to organize existing AWS best practice, NIST CSF functions, and the ML detection literature into a single reference architecture; it has not been empirically validated against live telemetry, and Section 7 and Section 9 return to this limitation and to the empirical work required to test it.

Layer 1 – Asset and Identity Governance (NIST Identify / Protect)

AWS Config supports ongoing asset discovery and classification; IAM Access Analyzer verifies enforcement of least-privilege IAM policies; multi-factor authentication is required for all human identities; and AWS KMS is used to encrypt sensitive data. This layer establishes the secure baseline against which anomalies can be meaningfully detected.

Layer 2 – Multi-Source Telemetry Collection (NIST Detect)

AWS Cloud Trail (management, data, and Insights events), Amazon Guard Duty, VPC Flow Logs, and AWS Config change notifications are enabled simultaneously, with all telemetry routed to a single, immutable storage location. S3 Object Lock should be configured to prevent deletion by an adversary who has obtained sufficient credentials, and Cloud Trail log-file validation must be enabled to detect tampering [25].

Layer 3 – ML-Driven Anomaly Detection (NIST Detect)

This layer deploys a tiered machine learning detection stack combining: (a) semi-supervised models trained on

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

historical incident data and supplemented with unlabeled operational telemetry, for detecting known-variant APT patterns; (b) unsupervised models — auto encoders or clustering algorithms — for detecting novel, previously unseen APT TTPs across Cloud Trail and Guard Duty telemetry; and (c) rule-based SQL detection scripts in Amazon Athena for rapid identification of high-confidence APT indicators, such as the creation of administrative IAM users, the suspension of Cloud Trail, or the sharing of snapshots with unrecognized accounts [14], [15]. This multi-paradigm stack is intended to compensate for the individual shortcomings of each strategy discussed in Section 3.4.

Layer 4 – Incident Response and Recovery (NIST Respond / Recover)

Each identified APT tactic maps to a pre-defined response playbook drawing on Cloud Trail evidence-preservation procedures, AWS Systems Manager for automated remediation, and Amazon Detective for rapid investigation. To close the loop in line with the NIST CSF Recover function, post-incident review procedures should feed lessons learned back into ML model retraining and detection-rule refinement [24].

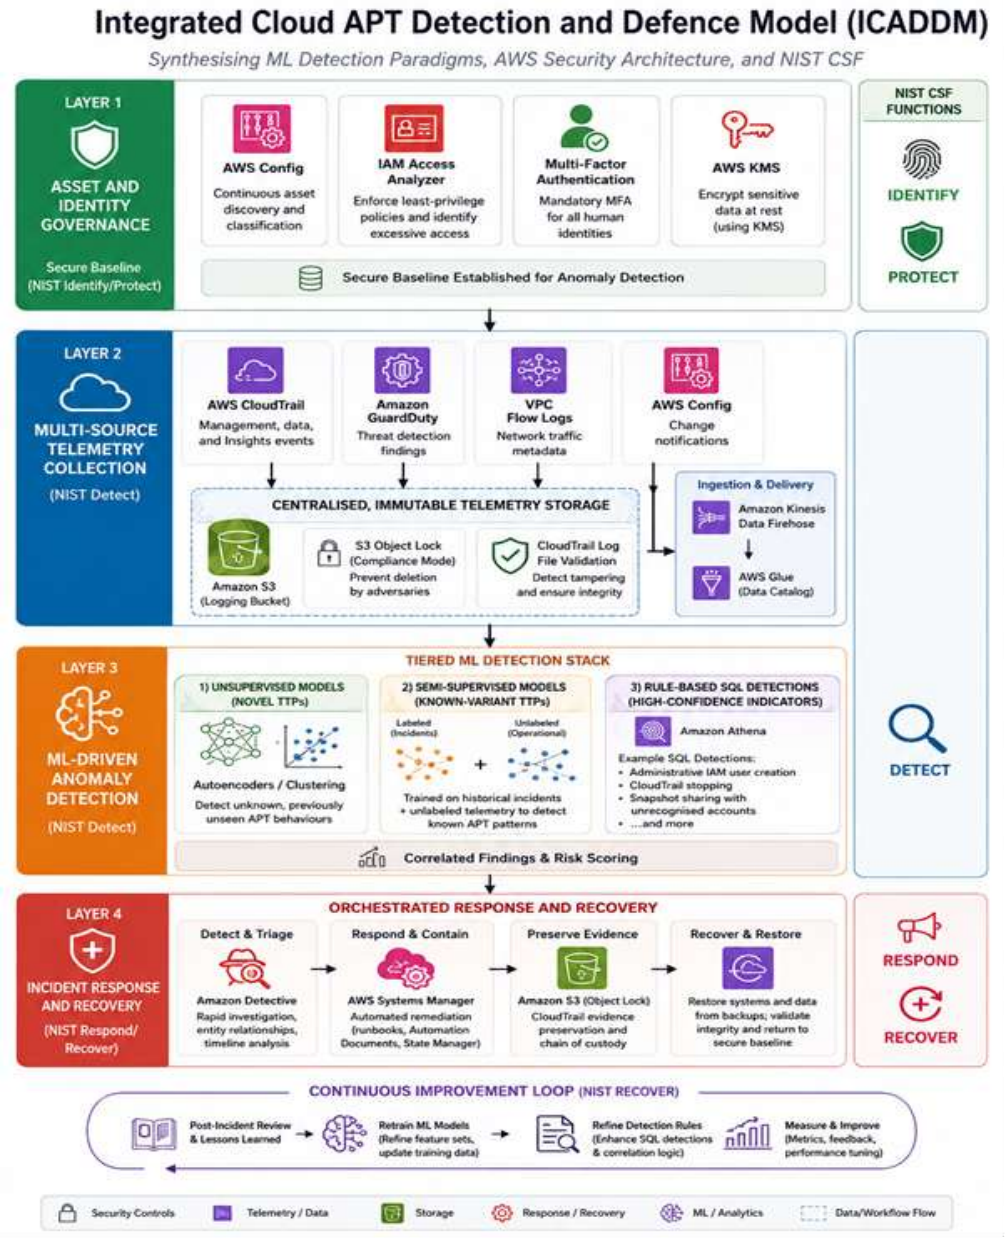

Figure 1. Proposed Integrated Cloud APT Detection and Defense Model (ICADDM)

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

DISCUSSION

This section discusses the comparative evaluation of ML paradigms and AWS security services for cloud APT detection set out in Sections 3 and 6.

The comparison in Table 3 suggests that no single machine learning paradigm is best suited to detecting cloud APTs on its own. Supervised learning offers high accuracy for established attack patterns, but its reliance on labeled training data and its consequent inability to recognize new APT TTPs is a significant constraint, given the dynamic and adaptive character of modern adversaries; because APT actors continually adapt their approach to avoid detection, reliance on supervised models alone is likely to be strategically insufficient [14].

Unsupervised learning addresses the novelty problem but introduces its own false-positive challenges, which can overwhelm security operations and cause alert fatigue, reducing operational effectiveness. Abhinav et al. [16] show how hybrid intrusion detection systems that combine anomaly-based and signature-based detection can offer more complete coverage by offsetting the drawbacks of each method individually. This finding directly supports the multi-paradigm detection stack proposed in the ICADDM, and suggests that the strategic combination of paradigms — rather than reliance on a single technique — is the more practical approach for cloud APT detection [16].

Deep learning models, particularly auto encoders and recurrent neural networks, represent the most capable paradigm for high-dimensional cloud telemetry analysis, though their deployment is constrained by processing-resource requirements and interpretability challenges [12], [14]. Relative to on-premises deployments, the resource constraint is somewhat less prohibitive in cloud environments, given the elastic scalability of compute resources; the associated cost of that elastic scaling, however, still needs to be weighed against detection benefit at the telemetry volumes discussed in Section 3.5.4. Security analysts continue to face a meaningful interpretability challenge, since understanding detection rationale is important for effective incident response; the combination of explainable AI techniques with deep learning models for cloud APT detection is accordingly an area that would benefit from further research.

The assessment of AWS logging and monitoring services in Table 4 indicates that each service addresses a different aspect of threat detection, and that their combined use produces a defense-in-depth architecture considerably stronger than any one service alone: AWS Cloud Trail provides the fundamental forensic record; Amazon Detective accelerates contextual analysis and investigation; AWS Config monitors configuration integrity; Amazon Inspector addresses vulnerability-level exposure; and Amazon Guard Duty provides automated, machine learning-based, real-time threat detection.

This review also identifies a structural weakness shared by organizations that rely heavily on Cloud Trail: an APT actor with sufficient IAM access can disable or delete the Cloud Trail, eliminating the primary source of forensic evidence. This defense-evasion technique underscores the need to treat Cloud Trail protection as a security control in its own right, rather than merely as an audit convenience. Hay and Nance [25] emphasize how central evidence integrity is to digital forensic investigation more generally — a principle that applies directly to cloud systems, where APT actors specifically target audit logs as part of their defense-evasion tactics.

Combining Cloud Trail's comprehensive event record with Amazon Guard Duty's machine learning-based threat detection produces a particularly effective detection pairing: Guard Duty generates real-time anomaly alerts using machine learning models, while Cloud Trail supplies the forensic context needed to investigate and verify those alerts. Amazon Detective further strengthens this combination by automatically correlating Guard Duty findings with Cloud Trail events, reducing mean time to investigation and allowing security analysts to focus on response rather than manual data assembly.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

In summary, the NIST CSF Detect function provides the governance mandate for implementing the multi-source telemetry collection and ML-based detection capabilities described in this paper, and is therefore central to any organizational program for cloud APT defense.

RECOMMENDATIONS

Based on the review presented in this paper, the following recommendations are offered to organizations seeking to strengthen their cloud APT detection and defense posture.

Implement a multi-paradigm machine learning detection stack.

Organizations should not rely on a single ML detection paradigm. The tiered approach underlying the ICADDM — combining unsupervised models for novel APT TTPs, semi-supervised models for known-variant patterns, and rule-based SQL scripts for high-confidence indicators — offers broader detection coverage. Amazon Sage Maker can be used to build and manage custom ML models at scale against Cloud Trail and Guard Duty telemetry [3].

Enable and safeguard all AWS telemetry sources concurrently.

Amazon Guard Duty, VPC Flow Logs, Cloud Trail (management, data, and Insights events), and AWS Config should all be enabled together. To prevent APT actors from destroying evidence, S3 Object Lock should be configured on the Cloud Trail storage bucket, and Cloud Trail log-file validation should remain active at all times [21], [25].

Invest in the development of cloud-specific APT datasets.

The scarcity of high-quality, labeled cloud APT datasets constrains ML model development. To generate the training data required for effective supervised and semi-supervised model deployment, organizations should participate in threat-intelligence sharing programs and adopt semi-synthetic dataset-generation techniques, as demonstrated by Myneni et al. [17].

Combine ML detection with human threat hunting.

Automated ML detection works best alongside human threat-hunting programs that actively search for APT indicators that automated models miss. Security analysts should routinely examine Guard Duty and Cloud Trail data using Amazon Detective to identify patterns that may indicate APT activity below the threshold of automated detection [10].

Account for adversarial ML risk when deploying ML-based APT detection.

Organizations should consider the possibility of adversarial machine learning — the intentional manipulation of detection models by capable APT attackers. An ML-operations program should incorporate adversarial-robustness testing, ensemble techniques that combine different model types, and regular retraining of models using recent telemetry [13].

CONCLUSION

This paper has reviewed AWS security architecture and machine learning-driven anomaly detection methods in the context of Advanced Persistent Threat mitigation in cloud environments. The review of ML paradigms — supervised, unsupervised, semi-supervised, and deep learning — shows that each carries distinct advantages and disadvantages, and that a multi-paradigm strategy compensating for each paradigm's individual shortcomings is

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

likely necessary for effective cloud APT detection. Unsupervised and semi-supervised learning appear most broadly applicable to cloud APT scenarios because they can identify new TTPs without depending on fully labeled training datasets.

AWS offers a wide range of security services — IAM, encryption, Cloud Trail, Guard Duty, Config, Detective, and Inspector among them — but the review of AWS security architecture in this paper indicates that these services must be deployed concurrently and in an integrated manner to achieve meaningful defense-in-depth against APT campaigns. On this basis, the paper has proposed the ICADDM as a conceptual architecture combining asset and identity governance, multi-source telemetry collection, ML-driven anomaly detection, and incident response and recovery, mapped to the NIST CSF and to representative MITRE ATT&CK Cloud tactics.

LIMITATIONS AND FUTURE RESEARCH

This work is a narrative and conceptual review rather than an empirical study, and the ICADDM proposed in Section 6 has not been validated against live AWS Cloud Trail data from real APT campaigns; such validation would substantially strengthen the evidentiary basis for the model and is the most important direction for future work. Future studies should evaluate the adversarial robustness of ML detection models in cloud environments, empirically compare ML paradigms against cloud-specific APT telemetry, and examine the effectiveness of the ICADDM longitudinally across different organizational cloud architectures. Given the practical importance of detection interpretability for effective analyst response, combining explainable AI techniques with deep learning models for cloud APT detection is a particularly pressing research need. Future work should also empirically test the cost and latency trade-offs of scaling ML-based detection across the multi-million-event daily telemetry volumes characteristic of large multi-account AWS environments, and should extend the illustrative MITRE ATT&CK mapping in Table 6 into a systematically validated technique-level coverage assessment.

REFERENCES

1. Brewer, R. (2014). Advanced persistent threats: Minimising the damage. Network Security, 2014(4), 5–9. https://doi.org/10.1016/S1353-4858(14)70040-6
2. Khaleefa, E. J., & Abdulah, D. A. (2022). Concept and difficulties of advanced persistent threats (APT): Survey. International Journal of Nonlinear Analysis and Applications. http://dx.doi.org/10.22075/ijnaa.2022.6230
3. Sharma, A., Gupta, B. B., Singh, A. K., & Saraswat, V. K. (2023). Advanced persistent threats (APT): Evolution, anatomy, attribution and countermeasures. Journal of Ambient Intelligence and Humanized Computing, 1–27.
4. Okoli, U., Obi, O., Adewusi, A., & Abrahams, T. (2024). Machine learning in cybersecurity: A review of threat detection and defense mechanisms. International Journal of Cyber-Security and Digital Forensics, 9, 147–154. https://doi.org/10.17781/P002677
5. Teichmann, F., Boticiu, S. R., & Sergi, B. S. (2023). The evolution of ransomware attacks in light of recent cyber threats. International Cybersecurity Law Review, 4, 259–280.
6. National Institute of Standards and Technology. (2011). Managing information security risk: Organization, mission, and information system view (NIST Special Publication 800-39). U.S. Department of Commerce. https://doi.org/10.6028/NIST.SP.800-39
7. Ghafir, I., & Prenosil, V. (2015). Advanced persistent threat and spear phishing emails. In M. Hrubý (Ed.), Proceedings of the International Conference Distance Learning, Simulation and Communication 'DLSC 2015' (pp. 34–41). University of Defence.
8. Alshaikh, A., Alanesi, M., Yang, D., & Alshaikh, A. (2023). Advanced techniques for cyber threat intelligence-based APT detection and mitigation in cloud environments. In P. Loskot & S. Niu (Eds.), Proceedings of the International Conference on Cyber Security, Artificial Intelligence, and Digital Economy (CSAIDE 2023) (Vol. 12718, Article 127180M). SPIE. https://doi.org/10.1117/12.2681627

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue VII, July 2026

9. Khan, S., Kabanov, I., Hua, Y., & Madnick, S. (2022). A systematic analysis of the Capital One data breach: Critical lessons learned. ACM Transactions on Privacy and Security, 26(1), Article 3. https://doi.org/10.1145/3546068
10. Cybersecurity and Infrastructure Security Agency, Federal Bureau of Investigation, Royal Canadian Mounted Police, Australian Cyber Security Centre, Australian Federal Police, Canadian Centre for Cyber Security, & National Cyber Security Centre (UK). (2025). Scattered Spider (Cybersecurity Advisory AA23-320A, updated). Cybersecurity and Infrastructure Security Agency. https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-320a
11. MITRE Corporation. (2024). Cloud matrix. MITRE ATT&CK. https://attack.mitre.org/matrices/enterprise/cloud/
12. Mandiant (Google Cloud). (2025). M-Trends 2025 special report. Google Cloud Security. https://services.google.com/fh/files/misc/m-trends-2025-en.pdf
13. Li, Z., Ge, Y., Guo, J., Chen, M., & Wang, J. (2022). Security threat model under internet of things using deep learning and edge analysis of cyberspace governance. International Journal of Systems Assurance Engineering and Management, 13, 1164–1176.
14. Saabith, A. S., Vinothraj, T., Fareez, M. M. M., & Marzook, M. M. (2023). A survey of machine learning techniques for anomaly detection in cybersecurity. International Journal of Research in Engineering and Science (IJRES), 11(10), 183–193.
15. Neuschmied, H., Winter, M., Stojanović, B., Hofer-Schmitz, K., Boegl, U., & Kleb, U. (2022). APT-attack detection based on multi-stage autoencoders. Applied Sciences, 12(13), 6816. https://doi.org/10.3390/app12136816
16. Abhinav, R., Raghav, K. N., Reddy, S. S., Koushik, P. S., Thangavel, S. K., & Srinivasan, K. (2023). A cloud-based intrusion detection system for advanced threat detection and prevention using machine learning techniques. 2023 14th International Conference on Computing Communication and Networking Technologies (ICCCNT), 1–8.
17. Myneni, S., Jha, K., Deng, Y., Chowdhary, A., Pisharody, S., & Huang, D. (2023). Unraveled — a semi-synthetic dataset for advanced persistent threats. Computer Networks, 227, 109688. https://doi.org/10.1016/j.comnet.2023.109688
18. Amazon Web Services. (2024). Logging and events. AWS Security Incident Response Guide. https://docs.aws.amazon.com/whitepapers/latest/aws-security-incident-response-guide/logging-and-events.html
19. Singh, T. (2021, November). The effect of Amazon Web Services (AWS) on cloud-computing. International Journal of Engineering Research & Technology (IJERT), 10(11).
20. Amazon Web Services. (2024). AWS Key Management Service (KMS). https://aws.amazon.com/kms/
21. Amazon Web Services. (2025). AWS CloudTrail user guide. https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-user-guide.html
22. Amazon Web Services. (2025). What is Amazon Security Lake? Amazon Security Lake user guide. https://docs.aws.amazon.com/security-lake/latest/userguide/what-is-security-lake.html
23. Amazon Web Services. (2022, January 5). Announcing AWS CloudTrail Lake, a managed audit and security lake. AWS What's New. https://aws.amazon.com/about-aws/whats-new/2022/01/aws-cloudtrail-lake-audit-security
24. National Institute of Standards and Technology. (2018). Framework for improving critical infrastructure cybersecurity, Version 1.1 (NIST CSWP 04162018). U.S. Department of Commerce. https://doi.org/10.6028/NIST.CSWP.04162018
25. Hay, B., & Nance, K. (2008). Forensics examination of volatile system data using virtual introspection. ACM SIGOPS Operating Systems Review, 42(3), 74–82. https://doi.org/10.1145/1368506.1368517