International Journal of Advanced Multidisciplinary Research and Educational Development
Volume 1, Issue 4 | November - December 2025 | www.ijamred.com
ISSN: **3107-6513**

Hybrid Log-Based Intrusion Detection for Control-Plane API Attacks in Cloud Environments

Samip Sanas[1], Amit Patil[2], Sandhya Kaprawan[3]

[1]M.S.(Cybersecurity),[2]M.S.(Cybersecurity, [3]Assistant Professor

[1,2,3]University Department of Information Technology, University of Mumbai, Kalina, Maharashtra, India

[1]samip2613@gmail.com, [2]amitp58476@gmail.com, [3]sandhya.kaprawan@udit.mu.ac.in

***Abstract*—Cloud control-plane APIs manage critical cloud resources and therefore represent a high-impact attack surface when misused through stolen credentials or policy abuse. Traditional intrusion detection systems are largely ineffective against such attacks because they rely on network signatures or isolated events rather than behavioral context. This paper presents a design-level analytical study of a hybrid detection framework for control-plane API attacks using log-based analytics. The proposed approach combines rule-based SIEM detection with a BiGRU-CNN model to conceptually capture both known misuse patterns and temporal anomalies from simulated AWS CloudTrail logs. The paper analytically examines expected detection behavior and design trade-offs, while practical implementation and empirical validation are deferred to future work.**

Keywords—Cloud Control-Plane Security, AWS CloudTrail Logs, Hybrid Intrusion Detection, Log-Based Analytics, SIEM and Machine Learning.

I. INTRODUCTION

Cloud computing platforms rely heavily on control-plane Application Programming Interfaces (APIs) to manage critical operations such as identity and access management, resource provisioning, policy enforcement, and service configuration. These APIs operate with high privileges and centralized authority, meaning that any misuse or compromise can result in severe security consequences, including privilege escalation, data exposure, and full cloud account takeover. Recent cloud security reports and studies highlight that control-plane misuse and identity-based attacks are among the most impactful threats in modern cloud environments [16], [17].

Traditional intrusion detection systems (IDS) were primarily designed for on-premise or network-centric environments and focus on packet inspection, signature matching, or isolated event analysis. Such approaches are often ineffective in cloud environments, where attackers commonly abuse valid credentials and invoke legitimate APIs in malicious sequences. Since control-plane attacks frequently resemble normal administrative activity, signature-based and perimeter-focused IDS solutions struggle to distinguish malicious behavior from legitimate cloud operations [1], [13].

To address these limitations, recent research has explored machine learning-based intrusion detection techniques that analyze logs and behavioral patterns rather than individual events. Studies show that machine learning and deep learning models can improve detection of complex attack behaviors, but they often introduce challenges related to interpretability, detection latency, and operational complexity when used as standalone solutions [6], [14]. Conversely, rule-based Security Information and Event Management (SIEM) systems provide fast and explainable detection but lack resilience against evolving and stealthy attack strategies [4].

This paper proposes a hybrid detection framework that combines rule-based SIEM analytics with deep learning-based behavioral modeling to conceptually address control-plane API attacks. The approach is designed around log-based analytics using simulated AWS CloudTrail events, enabling analysis of both known misuse patterns and anomalous temporal behavior [11], [18]. The scope of this paper is limited to design and analytical evaluation of the proposed framework. Practical implementation, model training, and real-world validation are intentionally deferred to a subsequent phase of research. Unlike prior intrusion detection studies, this work focuses exclusively on control-plane API misuse and analytically integrates SIEM-based rules with a BiGRU-CNN behavioral model using cloud audit logs.

II. RELATED WORK

Research on intrusion detection in cloud environments has increasingly shifted from traditional network-centric approaches to log-based and machine learning-driven techniques. Early studies on cloud intrusion detection systems (IDS) primarily focused on adapting signature-based and rule-driven mechanisms to virtualized infrastructures. While these methods offer low latency and interpretability, they struggle to detect sophisticated attacks that exploit legitimate credentials or cloud-native APIs, particularly in dynamic and multi-tenant environments [1], [13].

Machine learning and deep learning techniques have been widely explored to overcome the limitations of static detection systems. Several studies demonstrate that classifiers such as decision trees, support vector machines, and ensemble methods

 ISSN: 3107-6513

improve detection accuracy by learning complex patterns from cloud data [6], [14]. More recent work applies deep learning architectures, including CNNs and recurrent models, to capture spatial and temporal characteristics of attack behavior, showing improved performance against previously unseen threats [3], [15]. However, these approaches are highly dependent on training data quality and often lack transparency, making operational deployment challenging.

Log-based intrusion detection has gained attention due to the availability of structured audit logs generated by cloud platforms. Research indicates that cloud logs such as AWS CloudTrail provide valuable contextual information, including user identity, API call sequences, and resource metadata, which are well-suited for behavioral analysis [11], [18]. Advanced approaches incorporate correlation engines and language-model-based analysis to enhance detection coverage, though such techniques introduce higher computational overhead and complexity [9], [12].

Security Information and Event Management (SIEM) systems remain a core component of enterprise cloud security due to their ability to correlate events and apply predefined detection rules. Studies on SIEM rule optimization highlight their effectiveness in identifying known attack patterns but also emphasize their inability to adapt to evolving threats without manual updates [4]. Recent research suggests that hybrid detection frameworks combining rule-based SIEM logic with machine learning models provide a balanced solution, achieving improved detection accuracy while maintaining explainability and acceptable response times [10], [12].

Despite these advancements, existing literature reveals a gap in detection strategies specifically targeting cloud control-plane API abuse. Most prior work concentrates on network traffic or application-layer attacks, with limited focus on the high-privilege operations performed through control-plane APIs [17]. This gap motivates the proposed hybrid, log-based analytical framework, which is designed to conceptually address control-plane API attacks using simulated cloud audit logs.

Approach Type	Data Source	Technique Used	Key Limitation
ML-based IDS	Cloud audit logs	Traditional ML classifiers	Lacks control-plane–specific focus
Deep learning IDS	Cloud traffic & logs	Hybrid CNN-LSTM	High computational complexity
AI-based IDS	Cloud infrastructure logs	Machine learning models	Limited interpretability
Comparative IDS study	Network & cloud data	Multiple ML algorithms	No API misuse analysis
Deep learning IDS	Cloud services	Deep learning models	High false-positive rate

Table 1. Comparison of Existing Cloud Intrusion Detection Approaches

III. METHODOLOGY

This chapter describes the proposed methodology for analytically examining a hybrid detection framework aimed at identifying control-plane API attacks in cloud environments. The methodology is intentionally design-oriented and does not include implementation-specific details such as model training parameters or deployment configurations. Instead, it focuses on the logical detection pipeline and interaction between system components.

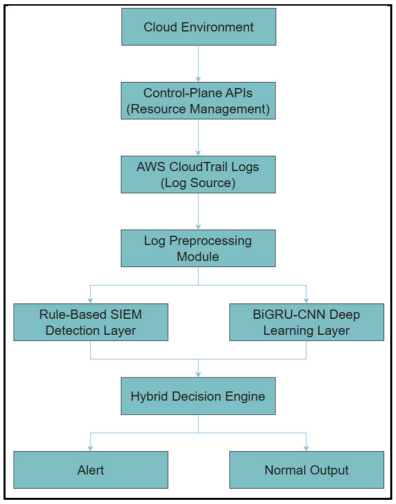

Figure 1. Architecture of the Proposed Hybrid Control-Plane Intrusion Detection Framework

3.1 Log Source and Simulation Strategy

The proposed framework is designed to operate on cloud audit logs generated by control-plane activities. For this study, simulated AWS CloudTrail logs are considered as the primary data source. CloudTrail logs provide structured records of management-level API calls, including user identity, event source, action type, timestamp, and request context [18]. Simulation is adopted to avoid direct dependency on live cloud environments and to ensure controlled representation of both normal and malicious behaviors, as commonly practiced in analytical cloud security studies [11].

3.2 Rule-Based Detection Logic

The first layer of the proposed framework consists of a rule-based detection mechanism inspired by traditional SIEM systems. This layer evaluates incoming log events against predefined rules that represent known misuse patterns, such as repeated failed authentication attempts, anomalous privilege changes, or suspicious resource modification sequences. Rule-based detection offers low latency and high interpretability, allowing immediate identification of well-understood attack behaviors [4]. However, its reliance on static rules limits its ability to detect novel or evolving attack strategies.

3.3 BiGRU-CNN Detection Logic

To complement rule-based detection, the second layer employs a conceptual deep learning model based on a hybrid Bidirectional Gated Recurrent Unit (BiGRU) and Convolutional Neural Network (CNN) architecture. In this

 ISSN: **3107-6513**

design, the CNN component is intended to extract local feature patterns from encoded log events, while the BiGRU component models bidirectional temporal dependencies across sequences of API calls. This combination is well-suited for capturing behavioral anomalies and multi-stage attack patterns that may not trigger explicit rules [3], [15]. The model operates on event sequences rather than isolated logs, enabling behavioral context analysis.

3.4 Hybrid Decision Flow

The final detection outcome is determined through a conceptual hybrid decision flow that integrates outputs from both the rule-based and deep learning layers. Events identified with high confidence by the rule-based layer are prioritized for immediate alerting, while outputs from the BiGRU-CNN layer are used to assess anomalous behavior patterns over time. By combining deterministic rules with adaptive learning-based inference, the hybrid approach aims to balance interpretability, adaptability, and detection coverage, as suggested in prior hybrid intrusion detection research [10], [12].

This methodology provides a structured analytical framework for studying control-plane API attack detection using log-based analytics. The practical implementation, model training, and empirical evaluation of this framework are intentionally deferred to future work.

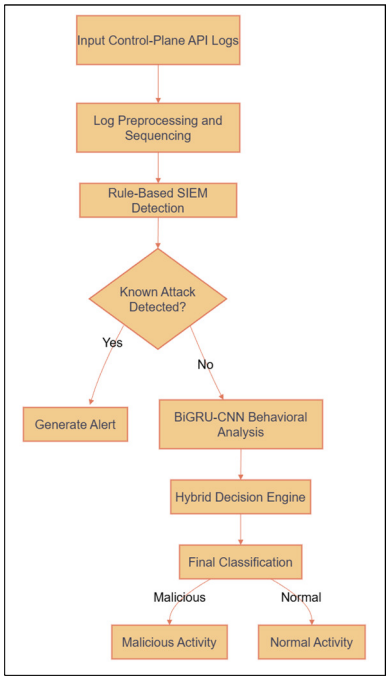

Figure 2. Hybrid Detection Workflow for Control-Plane API Attacks

IV. ASSUMED DATA

Reliable analysis of intrusion detection mechanisms for cloud environments depends strongly on the nature and realism of the data under consideration. Since this study focuses on a design-level and analytical evaluation, the proposed framework is examined using assumed and simulated control-plane API log data rather than logs collected from a live cloud deployment. This approach enables controlled analysis while avoiding ethical, security, and access constraints commonly associated with real cloud environments [11], [18].

4.1 Rationale for Using Simulated Control-Plane Logs

Simulated AWS CloudTrail logs are assumed as the primary data source because CloudTrail provides detailed, structured records of control-plane API activity, including user identity, API action, source context, and timestamps [18]. Using simulated data allows the modeling of both benign and malicious behaviors in a controlled manner, ensuring coverage of attack scenarios that may be rare or difficult to observe in production environments. Similar simulation-based assumptions are commonly adopted in analytical cloud security research to study detection behavior without claiming real-world deployment [1], [6].

4.2 Assumed Attack Behaviors

The assumed malicious behaviors are derived from documented cloud attack techniques and threat reports. These include misuse of valid credentials, anomalous privilege escalation through policy modification, unusual resource provisioning patterns, and abnormal sequences of control-plane API calls [16], [17]. Such behaviors are assumed to manifest as deviations in event frequency, role usage, API call ordering, or access context rather than as isolated malicious requests. These assumptions reflect the nature of control-plane attacks, which often rely on stealth and persistence rather than explicit exploit payloads.

4.3 Normal Versus Malicious Control-Plane Activity

Normal control-plane activity is assumed to consist of predictable and role-consistent API usage patterns, such as routine resource management actions performed by authorized identities during standard operational timeframes. In contrast, malicious activity is assumed to exhibit behavioral irregularities, including unexpected API calls by low-privilege identities, rapid changes in access policies, or uncommon sequences of management operations. The distinction between normal and malicious behavior is therefore treated as a behavioral classification problem rather than a signature-matching task, consistent with prior log-based intrusion detection studies [9], [12].

4.4 Design Realism and Practical Relevance

The design assumptions made in this paper are grounded in publicly available cloud documentation and established security research. By aligning simulated log structures and assumed behaviors with AWS CloudTrail semantics and recognized threat patterns, the proposed framework maintains practical relevance despite the absence of real deployment data [16], [18]. These design considerations ensure that the analytical insights derived from this study can be directly applied to future implementation and validation efforts.

 ISSN: **3107-6513**

This chapter establishes a transparent foundation for the analytical discussion that follows, clarifying the scope and limitations of the assumed data while ensuring that the proposed detection framework remains realistic and applicable to real-world cloud security scenarios.

V. ANALYTICAL DISCUSSION

This section presents an analytical discussion of the proposed hybrid detection framework based on the assumed data characteristics and design considerations outlined in the previous chapter. The discussion focuses on expected detection behavior, comparative strengths with respect to existing approaches, and practical design trade-offs, without claiming empirical implementation or measured performance.

5.1 Expected Detection Behavior

Based on the proposed design, the rule-based detection layer is expected to effectively identify well-known control-plane misuse patterns, such as repeated authorization failures, suspicious policy modifications, and abnormal resource configuration changes. Due to its deterministic nature, this layer provides fast and explainable alerts, making it suitable for identifying high-confidence misuse scenarios [4].

The BiGRU-CNN detection layer is analytically expected to complement rule-based detection by identifying behavioral anomalies that emerge over sequences of API calls. By modeling temporal dependencies and contextual patterns, this layer can capture stealthy or multi-stage control-plane attacks that may not trigger explicit rules, such as gradual privilege escalation or low-frequency anomalous operations [3], [15]. The hybrid decision flow is therefore expected to improve overall detection coverage by combining immediate rule-based responses with adaptive behavioral analysis.

5.2 Comparison with Prior Studies

Prior research indicates that rule-based intrusion detection systems perform well against known attack patterns but exhibit limited adaptability to evolving threats, particularly in cloud-native environments [1], [13]. Conversely, machine learning-based approaches demonstrate improved generalization capabilities but often suffer from higher detection latency and reduced interpretability when deployed independently [6], [14].

The proposed hybrid framework conceptually aligns with studies advocating combined detection strategies, where rule-based systems provide baseline security while learning-based models enhance detection of novel behaviors [10], [12]. Unlike many existing approaches that focus on network traffic or application-layer logs, this framework emphasizes control-plane API logs, addressing a gap identified in cloud security literature [17].

Detection Strategy	Adaptability	Interpretability	Control-Plane Suitability
Rule-Based IDS	Low	High	Medium
ML-Based IDS	High	Low	Medium
Proposed Hybrid IDS	High	High	High

Table 2. Comparison of Rule-Based, ML-Based, and Hybrid Detection Strategies

5.3 Strengths of the Hybrid Design

The primary analytical strength of the proposed framework lies in its balanced design. Rule-based detection ensures low-latency and transparent alerting, which is critical for operational cloud security monitoring. The BiGRU-CNN layer adds adaptability by learning temporal patterns in control-plane activity, enabling detection of complex and distributed attack behaviors. Together, these components are analytically expected to reduce false positives while improving detection of subtle anomalies compared to standalone approaches [9], [12].

Additionally, the use of structured CloudTrail logs allows the framework to operate without deep packet inspection, making it suitable for cloud environments where network-level visibility is limited or abstracted [11], [18].

5.4 Design Trade-offs and Limitations

Despite its analytical advantages, the proposed hybrid framework involves inherent trade-offs. The inclusion of a deep learning component introduces additional computational overhead and complexity compared to purely rule-based systems. While this overhead is analytically justified by improved detection capability, it may impact real-time performance in large-scale environments if not carefully optimized.

Furthermore, the analytical discussion presented in this paper is based on assumed and simulated data. As such, the actual effectiveness of the framework in real-world deployments depends on implementation details, data quality, and operational constraints, which are intentionally deferred to future work. These limitations are acknowledged to ensure transparency and to clearly delineate the scope of this analytical study

VI. CONCLUSION

Control-plane APIs form the backbone of cloud infrastructure management and, due to their high privilege level, represent a critical target for attackers. Traditional intrusion detection systems, which rely primarily on network signatures or isolated event analysis, are insufficient for identifying control-plane misuse that often leverages valid credentials and legitimate API calls. This paper presented a design-level and analytical study

of a hybrid detection framework aimed at addressing these limitations using log-based analytics.

The proposed framework conceptually integrates rule-based SIEM detection with a BiGRU-CNN–based behavioral analysis model to improve coverage against both known misuse patterns and subtle, time-distributed anomalies in control-plane activity. By operating on simulated AWS CloudTrail logs, the study analytically examined expected detection behavior, design strengths, and trade-offs without claiming empirical implementation or measured performance. The analysis indicates that a hybrid approach can balance interpretability, adaptability, and operational feasibility more effectively than standalone rule-based or machine learning-based solutions.

The scope of this paper is intentionally limited to design assumptions and analytical discussion. Practical implementation, model training, performance evaluation, and real-world validation using live or large-scale cloud datasets are deferred to future work. The analytical framework established in this study serves as a foundation for the subsequent implementation phase, which will focus on building, deploying, and empirically validating the proposed hybrid detection system in a cloud environment.

REFERENCES

[1] K. Mahendar and G. Shivakanth, "A survey of intrusion detection systems based on machine learning for cloud security," *SSRG International Journal of Electrical and Electronics Engineering*, vol. 12, no. 5, pp. 119–128, 2025.

[2] M. Sohail, "Intelligent threat detection and prevention in REST APIs using machine learning," *International Journal of Scientific Research Archive*, vol. 15, no. 2, pp. 012–027, 2025, doi: 10.30574/ijsra.2025.15.2.1281.

[3] M. M. Alshehri, S. A. Alshehri, and S. H. Alajmani, "Intrusion detection system using hybrid CNN-LSTM model in cloud environment," *Indonesian Journal of Electrical Engineering and Computer Science*, vol. 40, no. 2, pp. 840–849, 2025, doi: 10.11591/ijeecs.v40.i2.pp840.

[4] A. Shukla, P. A. Gandhi, Y. Elovici, and A. Shabtai, "RuleGenie: SIEM detection rule set optimization," *arXiv preprint*, arXiv:2505.06701, 2025.

[5] IRJIET Research Team, "Enhancing cloud security in AWS using AI-powered anomaly detection and predictive analytics," *International Research Journal of Innovative Engineering and Technology*, vol. 9, no. 3, 2025.

[6] M. M. Khan, "Developing AI-powered intrusion detection system for cloud infrastructure," *Journal of Artificial Intelligence, Machine Learning and Data Science*, vol. 2, no. 1, pp. 1074–1080, 2024, doi: 10.51219/JAIMLD/mohammed-mustafa-khan/255.

[7] A. Rathore and T. Sahu, "AI-based intrusion detection system in cloud computing," *International Journal of Innovative Research in Computer Science and Technology*, vol. 12, no. 1, pp. 45–51, 2024, doi: 10.55524/CSISTW.2024.12.1.8.

[8] V. K. Gandam and E. Aravind, "Enhancing cloud security: A novel intrusion detection system using deep learning algorithms," *International Journal of Computer Applications*, vol. 186, no. 44, pp. 36–42, 2024, doi: 10.5120/ijca2024924070.

[9] W. Guan *et al.*, "LogLLM: Log-based anomaly detection using large language models," *arXiv preprint*, arXiv:2411.08561, 2024.

[10] N. Wankhade and A. Khandare, "Optimization of deep generative intrusion detection system for cloud computing: Challenges and scope for improvements," *EAI Endorsed Transactions on Scalable Information Systems*, vol. 10, no. 6, 2023, doi: 10.4108/eetsis.3993.

[11] S. F. Javior, "Log-based intrusion detection system using machine learning," M.S. thesis, Dept. of Computing, National College of Ireland, Dublin, Ireland, 2023.

[12] M. Sheeraz *et al.*, "Revolutionizing SIEM security: An innovative correlation engine design for multi-layered attack detection," *Sensors*, vol. 24, no. 15, p. 4901, 2024, doi: 10.3390/s24154901.

[13] G. Rathod, V. Sabnis, and J. K. Jain, "Intrusion detection system (IDS) in cloud computing using machine learning algorithms: A comparative study," *Grenze International Journal of Engineering and Technology*, vol. 10, no. 1, pp. 550–563, 2022.

[14] I. Hidayat, M. Z. Ali, and Arshad, "Machine learning-based intrusion detection system: An experimental comparison," *Journal of Computational and Cognitive Engineering*, vol. 2, no. 2, pp. 88–97, 2022, doi: 10.47852/bonviewJCCE2202270.

[15] S. S. Chakravarthi *et al.*, "Deep learning-based intrusion detection in cloud services for resilience management," *Computer Modeling in Engineering & Sciences*, vol. 71, no. 3, pp. 5117–5133, 2022, doi: 10.32604/cmc.2022.022351.

[16] European Union Agency for Cybersecurity (ENISA), *ENISA Threat Landscape 2023*, ENISA, Oct. 2023.

[17] S. Rizvi, A. Kurtz, J. Pfeffer, and M. Rizvi, "Securing cloud control planes: Challenges and solutions," *IEEE Security & Privacy*, vol. 20, no. 4, pp. 58–66, Jul.–Aug. 2022, doi: 10.1109/MSEC.2022.3156789.

[18] Amazon Web Services, *AWS CloudTrail User Guide*, AWS Documentation.