HCAPO: Transformer-Based Adaptive Policy Orchestration for Hybrid-Cloud Security

Geonmin Kim
Dept. of Artificial Intelligence Convergence
Chonnam National University
Gwangju, Korea
geonminkim@jnu.ac.kr

Yejin Kim
Dept. of Software Engineering
Chonnam National University
Gwangju, Korea
ye031010@jnu.ac.kr

Taerim Kim
Dept. of Artificial Intelligence
Chonnam National University
Gwangju, Korea
ktr0706@jnu.ac.kr

Kyungbaek Kim
Dept. of Artificial Intelligence Convergence
Chonnam National University
Gwangju, Korea
kyungbaekkim@jnu.ac.kr

Abstract— **Hybrid cloud environments combine heterogeneous infrastructures, distributed workloads, and region-level policy inconsistencies, making traditional static or rule-based security mechanisms insufficient against evolving threats. Prior research has explored machine learning–based threat detection, log-driven policy generation, variance stabilization, and edge-level caching, but these approaches remain fragmented and lack an integrated pipeline capable of proactive, AI-driven security management. To address this gap, we propose HCAPO (Hybrid-Cloud Adaptive Policy Orchestration), an AI-augmented framework that unifies predictive threat modeling, dynamic policy synthesis, and dual-layer (edge–central) enforcement. HCAPO incorporates (i) a Transformer-enhanced time-series predictor for estimating attack likelihood from multi-field log sequences, (ii) a deep hybrid embedding model that fuses textual, temporal, and categorical log features, (iii) a stabilized policy generator using cross-selection to reduce feature variance, (iv) a Security Risk Score (SRS) that prioritizes inferred policies for edge deployment, and (v) an adaptive orchestration loop that continuously updates policies based on real-time feedback. Experiments using large-scale hybrid-cloud traffic logs demonstrate that HCAPO significantly improves predictive accuracy, detection performance, policy stability, and edge-level enforcement latency compared to ML-only or static policy baselines. These results highlight the feasibility of an integrated, AI-driven security pipeline and point toward fully autonomous, self-optimizing security management for hybrid and multi-cloud environments.**

Keywords—Cloud Security, Log Analysis, AI-Driven Security Automation, Adaptive Policy Generation, Threat Prediction

I. Introduction

Hybrid cloud architectures—comprising on-premise systems, private clouds, and multiple public cloud regions—have become the dominant deployment model for modern enterprises. These environments offer operational flexibility and scalability, yet they also introduce fragmented network boundaries, heterogeneous access control mechanisms, and region-dependent latency characteristics[1-3]. As a result, hybrid clouds inherently amplify the complexity of security management. Attackers exploit multi-layered interactions between API gateways, microservices, and cloud-native components, while security teams struggle to maintain consistent policies, detect emerging threats, and respond in real time[4-7].

Traditional perimeter-based defenses and rule-driven intrusion detection systems (IDS) fall short in such settings. They rely heavily on manually crafted rules, static threat signatures, or event-level anomaly scoring, all of which are insufficient for workloads that dynamically shift across hybrid infrastructures. To overcome these constraints, recent studies have focused on machine learning (ML) and log-centric analysis to automate threat detection and derive policies from historical attack patterns. Such approaches have demonstrated promising results in web attack classification, feature-based policy generation, and edge-level enforcement. However, they remain limited in three key aspects. First, prior ML-driven systems largely treat security events as independent samples, ignoring temporal dependencies, sequence-level correlations, and evolving attack behaviors. This hinders their ability to forecast threats before they manifest proactively. Second, existing policy generation frameworks rely on sparse feature vectors and static selection mechanisms, resulting in high variance and poor stability across multiple training cycles. Third, although edge policy caching has been explored to reduce evaluation latency, current systems lack a unified mechanism that links threat prediction, policy synthesis, and policy deployment into a coherent orchestration pipeline.

To address these gaps, this paper proposes HCAPO (Hybrid-Cloud Adaptive Policy Orchestration), a unified, AI-augmented security framework that integrates predictive modeling, dynamic policy generation, and multi-layer enforcement. HCAPO introduces deep sequence modeling based on Transformer-enhanced time-series predictors to capture temporal attack patterns, enabling proactive threat estimation. It further employs a hybrid deep embedding model that fuses textual, categorical, and temporal fields from multi-modal cloud logs, improving the semantic expressiveness of features used for policy synthesis. To stabilize generated policies, HCAPO incorporates a cross-selection mechanism that reduces feature variance and enhances consistency across retraining cycles. This design transitions hybrid cloud security from reactive detection to proactive, self-optimizing management.

II. Related Work

Research on hybrid-cloud security spans multiple domains, including cloud orchestration, risk-aware policy management, log-based anomaly detection, and AI-driven automation. This section reviews prior work across these areas and positions the proposed HCAPO framework within the evolving landscape of hybrid-cloud security research.

A. Hybrid and Multi-Cloud Orchestration

Hybrid- and multi-cloud architectures aim to combine the strengths of on-premise and public cloud infrastructures, but this integration introduces operational fragmentation. Sitaram et al. [1] examined interoperability and orchestration challenges across heterogeneous cloud providers, emphasizing the absence of standardized mechanisms for consistent policy execution. Mansouri et al. [2] implemented an automated hybrid-cloud environment to evaluate distributed database performance, revealing that workload placement, cross-cloud communication, and network topology significantly influence system behavior. Their findings underscore the need for adaptive orchestration capable of responding to dynamic performance conditions. At the service layer, Li et al. [3] conducted an empirical study analyzing cloud API reliability issues, identifying transient failures, inconsistent latency, and API-level anomalies across large-scale cloud systems. These issues become more pronounced in hybrid-cloud environments where API-driven interactions span diverse regions and infrastructures. While these works clarify the architectural challenges of hybrid-cloud systems, they focus primarily on orchestration and performance optimization rather than linking operational observations to automated security decision-making.

B. Cloud Security Threats in Hybrid Clouds

Cloud security research has identified an expanding set of vulnerabilities driven by distributed architectures. Mallisetty et al. [5] provided a comprehensive review of cloud security challenges, including data breaches, insecure APIs, misconfigurations, and cross-tenant threats. Raktate et al. [6] further categorized security issues into identity mismanagement, API abuse, data leakage, and lateral movement. In hybrid and multi-cloud environments, Mishra et al.[7] analyzed the elevated cyber risks caused by interoperability gaps, inconsistent policy enforcement, and divergent control mechanisms across platforms. Their study highlights the need for adaptable security mechanisms capable of operating across multiple trust boundaries and infrastructure layers. However, these works primarily focus on identifying threats rather than generating enforceable security policies or orchestrating defense mechanisms across distributed components. They provide important motivation but lack implementation pathways for automated policy operations.

C. AI-driven Security Automation

AI and automation have been increasingly applied to security automation and classifying malicious traffic[8-14]. Seth et al.[12] demonstrated how AI and generative automation can support multi-cloud security management, assisting with compliance validation, anomaly detection, and operational optimization. However, their work remains limited to event-level automation and does not generate or deploy dynamic enforcement policies. Declarative policy systems have also emerged as a means of standardizing cloud governance. Paul et al.[13] integrated Open Policy Agent (OPA) into AWS cloud environments to automate compliance checks, reducing manual overhead. Despite their practical success, these systems rely heavily on manually authored rules and do not incorporate machine-driven policy synthesis.

Edge-based enforcement has also received attention. Recent hybrid-cloud studies demonstrated that caching high-frequency or latency-sensitive rules at edge gateways can reduce evaluation latency and improve system responsiveness[14]. While such approaches validate the benefits of distributing enforcement across edge and central layers, they do not incorporate predictive modeling or adapt policies dynamically in response to evolving threats. Overall, AI-driven cloud security research has focused on detection and compliance but lacks end-to-end integration of prediction, policy generation, stability management, and adaptive deployment—an integration central to HCAPO.

D. Deep Learning for Log-Based Anomaly Detection

Log anomaly detection has become one of the most active areas in cloud security research, with deep learning offering significant improvements over rule-based or statistical methods[15-18]. Landauer et al. [16] surveyed deep-learning approaches including LSTM models, autoencoders, CNNs, and attention-based architectures, highlighting their strengths in modeling complex log semantics and temporal dependencies. Their review also noted that most existing systems treat logs as single-modal text and lack integration with operational policy engines. Transformer models have strengthened this field further. TransLog [17] proposed a unified Transformer-based framework that can transfer knowledge across log domains to enhance anomaly detection and reduce template dependency. More recently, ADALog [18] introduced a masked-language self-attention model for unsupervised log anomaly detection, allowing direct processing of raw logs without parsing.

Despite these advancements, existing deep-learning approaches remain detection-centric. None translates anomaly signals into enforceable policies, nor do they integrate with risk scoring or distributed enforcement systems. This separation between detection and operational decision-making forms a core gap that HCAPO aims to address.

E. Summary and Positioning of This Work

Across orchestration, cloud threat analysis, AI-driven automation, and deep log anomaly detection, prior research has made important contributions but remains disconnected. Existing work only identifies anomalies but does not synthesize or deploy policies, relies on limited or single-modal feature representations, lacks predictive threat modeling that drives proactive adaptation, does not offer risk-aware prioritization for policy deployment, and rarely coordinates enforcement across edge and central components.

The proposed HCAPO framework bridges these gaps by integrating deep sequence modeling, multi-modal log embedding, ML-driven policy synthesis, stability-enhanced selection, and risk-aware orchestration into a single adaptive system for hybrid-cloud environments.

III. System Architecture and Methodology

Hybrid-cloud environments combine on-premise infrastructures, public clouds, and private clouds into a unified operational ecosystem. While this integration enables scalability and flexibility, it also introduces significant complexity in maintaining consistent security across heterogeneous systems. Logs generated by these environments include diverse modalities—textual request paths, categorical method types, numerical latency values, and temporal metadata—making it challenging to reason about system behavior through any single analytic approach. To address these challenges, the proposed Hybrid-Cloud

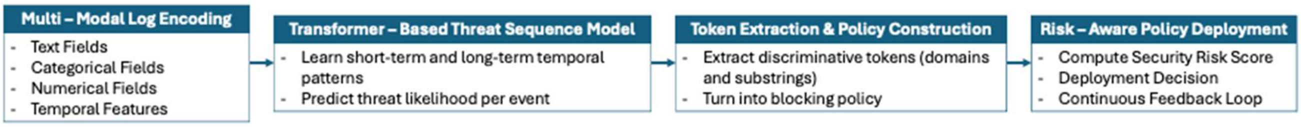
Multi – Modal Log Encoding
- Text Fields
- Categorical Fields
- Numerical Fields
- Temporal Features
Transformer – Based Threat Sequence Model
- Learn short-term and long-term temporal patterns
- Predict threat likelihood per event
Token Extraction & Policy Construction
- Extract discriminative tokens (domains and substrings)
- Turn into blocking policy
Risk – Aware Policy Deployment
- Compute Security Risk Score
- Deployment Decision
- Continuous Feedback Loop

Fig. 1. Transformer-Driven Threat Analysis and Risk-Aware Policy Generation Pipeline

Adaptive Policy Orchestration (HCAPO) framework integrates multi-modal representation learning, sequence-based threat prediction, token-guided rule synthesis, and risk-aware policy deployment within a unified pipeline. The architecture consists of three connected components: (1) Multi-Modal Log Understanding, (2) Transformer-Based Threat Prediction and Token Extraction, (3) Risk-Aware Policy Generation and Deployment. The integration of these components forms an end-to-end mechanism capable of continuously adapting security policies to evolving cloud conditions.

A. *Multi-Modal Log Understanding*

Logs arriving from gateways, microservices, virtual machines, and edge nodes vary substantially in their structure and meaning. To establish a common representation, each log entry is decomposed into four modalities: textual ($x_t^{(txt)}$), categorical ($x_t^{(cat)}$), numerical ($x_t^{(num)}$), and temporal ($x_t^{(time)}$). Each component is processed by a modality-specific encoder and later fused into a unified embedding. The fusion process can be abstractly expressed as:

$$h_t = f_{fusion}(h_t^{(txt)}, h_t^{(cat)}, h_t^{(num)}, h_t^{(time)}) \tag{1}$$

Here, h_t presents the semantically enriched representation of a log event, encapsulating structural, statistical, and temporal information. This unified representation allows downstream models to capture contextual relationships that are not visible when analyzing individual log components in isolation.

B. *Sequence-Based Threat Modeling*

Security threats in hybrid-cloud settings often emerge not from isolated events but from sequences of abnormal behaviors that accumulate over time. Thus, HCAPO employs a transformer-based sequence prediction model to analyze the temporal patterns embedded in $\{h_1, h_2, \dots, h_t\}$. The model learns both short-term deviations (e.g., sudden spikes in malicious URLs) and long-term behavioral drifts (e.g., gradual changes in access frequency or data exfiltration patterns). The model outputs an anomaly or threat likelihood score for each event and simultaneously generates interpretability artifacts such as attention heatmaps and salient token indicators. These interpretability signals become essential in the next stage, where the system synthesizes actionable policy rules.

C. *Token Extraction and Policy Construction*

Merely detecting anomalies is insufficient for practical security enforcement; the system must also transform model insights into deployable rules. To achieve this, HCAPO extracts discriminative log tokens—such as suspicious URL substrings, abnormal query parameters, or recurring malicious sequences—from the transformer's attention distribution and saliency signals.

Extracted tokens are mapped to predefined policy templates, which align with the operational semantics of edge gateways or OPA/Rego-based central evaluators. For instance, abnormal URL patterns may trigger deny rules, while repetitive suspicious parameters can be transformed into rate-limiting or header-validation policies. Before adoption, candidate rules undergo temporal stability analysis. Only rules that remain consistently relevant over multiple prediction cycles advance to final deployment, preventing temporary anomalies or noise from generating unstable or overly aggressive policies.

D. *Risk-Aware Policy Deployment*

Different layers of the hybrid-cloud environment vary in computational capacity and sensitivity to latency. Edge nodes enable low-latency enforcement but have limited processing capabilities, while centralized OPA clusters support complex logic at the cost of additional delay. HCAPO resolves this trade-off using a risk-based deployment mechanism.

Each candidate policy is evaluated by a composite risk score, integrating predicted threat likelihood (P_t), potential impact severity (S_i), and difficulty of detection (D_d):

$$S_{risk} = w_1 P_t + w_1 S_i + w_3 D_d \tag{2}$$

Policies exceeding a threshold τ are pushed to edge caches for immediate enforcement, whereas those below the threshold are routed to the centralized evaluator to benefit from richer context and stateful analysis. Formally:

$$Deploy(r) = \begin{cases} Edge, & S_{risk} \geq \tau \\ Central, & S_{risk} < \tau \end{cases}$$

This mechanism ensures that threats requiring immediate attention are mitigated at the earliest possible point, while broader, less urgent, or more computationally demanding rules are offloaded to central evaluation. A continuous feedback loop—capturing false positive events, cache hit ratios, and rule effectiveness—feeds back into the overall HCAPO pipeline, enabling long-term refinement and adaptive improvement of both the prediction model and policy generation logic.

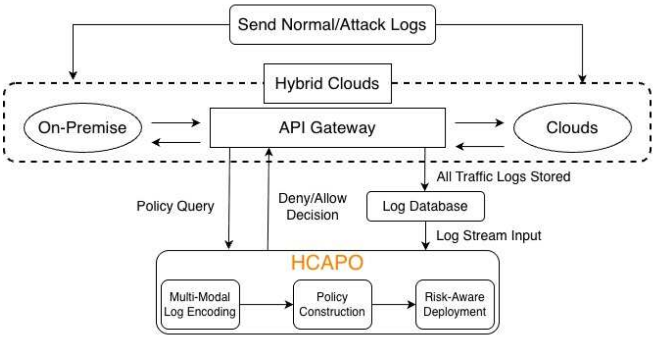

Fig. 2. System Architecture of the Proposed HCAPO Framework

IV. Experimental Setup

This section describes the dataset, preprocessing procedures, sequence construction method, model configuration, evaluation protocol, and implementation environment used to assess the HCAPO framework. The experiments were designed to emulate threat patterns observable in hybrid-cloud infrastructures, where DNS-based traffic plays a central role in both benign operations and malicious activity. The multi-modal characteristics of the dataset—comprising textual domain identifiers, categorical protocol attributes, numerical flow metrics, and temporal patterns—enable a comprehensive evaluation of the proposed architecture.

A. Dataset Description

The experiments rely on DNS flow logs derived from PCAP captures, representing four categories of traffic: benign activity and three malicious classes corresponding to malware, phishing, and spam behavior. Each flow contains structured information, including the queried domain name, associated DNS resource record types, protocol categories, flow duration, byte counts, packet numbers, and timestamps that reflect the chronological order of events.

The dataset consists of 81,698 malware flows, 43,348 phishing flows, and 30,371 spam flows, accompanied by a benign set used to model normal behavior. Because all flows originate from the same client source, they naturally form a temporally ordered sequence that is well suited for transformer-based modeling. This sequential characteristic also mirrors realistic operational settings, where suspicious DNS behavior often emerges gradually through repetitive or abnormal query patterns.

B. Data Preprocessing

Before training, all flows were chronologically sorted to preserve temporal dependencies. Domain names were decomposed into hierarchical components—subdomain, second-level domain, and top-level domain—and tokenized using a data-driven vocabulary to capture structural variations across domains. Numerical attributes such as flow duration and byte counts were normalized and clipped at extreme quantiles to reduce the influence of outliers, while categorical attributes were mapped to integer indices for embedding.

The logs were then reorganized into fixed-length sequences of 32 flows using a sliding-window method with partial overlap. Each sequence was assigned a label according to the dominant class contained within the window, enabling the model to learn behavioral patterns rather than isolated anomalies. This approach reflects practical threat-detection scenarios, in which the maliciousness of a flow becomes evident only when contextualized within a broader temporal pattern. To evaluate the predictive capability of the model under realistic conditions, the dataset was partitioned according to temporal order. The earliest 70% of the flows were used for training, the subsequent 10% for validation, and the remaining 20% for testing. This time-based split prevents information leakage and enforces the operational constraint that future behavior must be predicted solely from past observations. Such a protocol aligns with real-world hybrid-cloud environments, where evolving threat patterns require continual adaptation to newly emerging behaviors.

C. Model Configuration

The predictive component of HCAPO is a multimodal transformer encoder composed of four layers, a hidden dimension of 256, a feed-forward size of 512, and eight attention heads. Separate embedding modules encode textual domain tokens, categorical protocol attributes, numerical metadata, and temporal features, which are then fused into a unified representation. The model was trained using the AdamW optimizer with a batch size of 128. No additional baseline models were employed in the final evaluation, as the goal of this study is to assess the end-to-end behavior of the HCAPO framework rather than to compare model families. The evaluation therefore focuses on the internal components of HCAPO itself: sequence classification accuracy, token extraction behavior, SRS-based threat quantification, and adaptive policy deployment outcomes.

D. Evaluation Metrics

Model performance was assessed using accuracy, precision, recall, and F1-score computed on the final test set. Because the harm associated with misclassifying malicious traffic is asymmetric, special emphasis was placed on recall for the malware, phishing, and spam classes. In addition to classification metrics, the study evaluates token-level extraction quality and the behavior of the Security Risk Score (SRS), which ranks domain tokens according to their inferred maliciousness. Finally, the adaptive policy deployment mechanism was evaluated through a simulation that measures how policies are distributed between edge and central enforcement layers, how their SRS values differ, and how frequently the deployed rules are triggered by test-time traffic. This simulation reflects the operational objective of placing high-risk policies at the edge while assigning lower-risk rules to the central controller.

V. Experimental Results

This section presents the empirical findings obtained from evaluating the HCAPO framework. The results assess three key aspects of system behavior: (1) the accuracy of multi-class traffic classification based on DNS flow sequences, (2) the reliability of token extraction and risk quantification for policy generation, and (3) the behavior of the risk-aware policy deployment mechanism, which determines how learned policies are distributed across edge and central enforcement layers. The analysis is organized according to these three perspectives and is supported by the figures presented in this section.

A. Multi-Class Classification Model Performance

The classification performance of the transformer-based model was assessed using the Confusion Matrix shown in Fig. 3. The matrix demonstrates that the model achieves approximately 85–86% correct classification across all four categories, as indicated by the diagonal proportions derived from the raw counts. For benign traffic, about 86% of samples were correctly identified as benign, with roughly 14% misclassified into malicious categories. Malware traffic shows a similar pattern, with approximately 85% correctly detected and the remaining 15% distributed across phishing and spam labels. Phishing and spam traffic exhibit comparable accuracy levels, again with correct classification rates near 85%, and misclassification rates of 4–6% into other malicious classes. These error patterns are expected, as phishing and spam domains often share syntactic structures with malware-related domains, leading to natural ambiguity in DNS-based prediction.

Despite these similarities, the model maintains strong overall separation between benign and malicious behavior. The proportion of benign traffic misclassified as malicious remains relatively low, which is essential for minimizing operational false positives in hybrid-cloud deployments. Overall, the confusion matrix confirms that the model reliably captures temporal and lexical signals necessary for distinguishing diverse malicious behaviors.

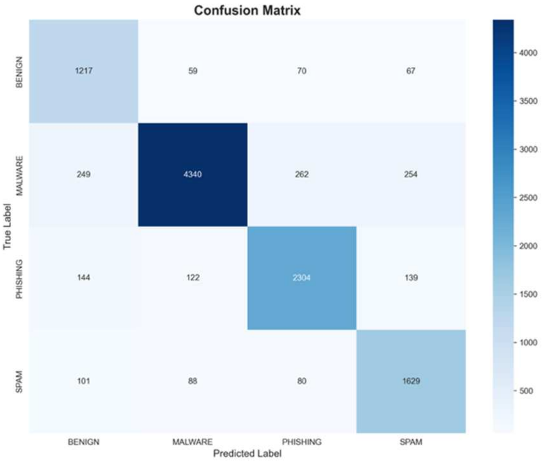

Fig. 3. Confusion Matrix (Raw Counts)

B. Token Extraction and Security Risk Scoring

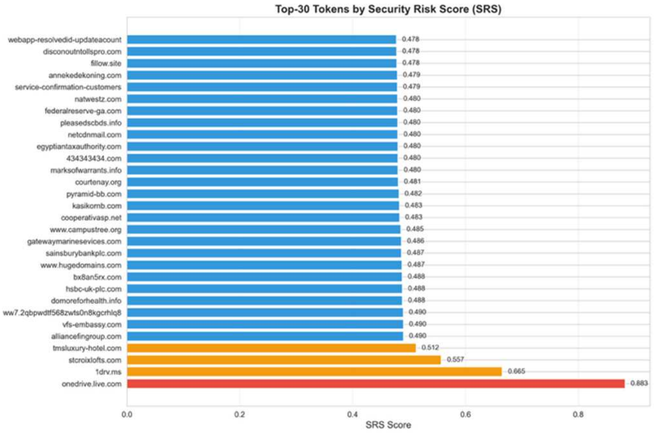

Fig. 4. Top-30 Tokens by Security Risk Score

Token extraction was performed on sequences identified as malicious, and each extracted token was assigned a Security Risk Score (SRS). Fig. 4 illustrates the top-ranked tokens, many of which correspond to domains frequently associated with credential harvesting, financial fraud, or impersonation campaigns. The top-scoring tokens exhibit SRS values concentrated near the upper end of the scoring range, indicating strong statistical correlation with malicious activity.

Many high-SRS tokens originate from syntactically unusual or irregular domain patterns, suggesting that the transformer encoder is effectively isolating structural anomalies. These findings validate the token extraction mechanism as an effective means for identifying fine-grained behavioral indicators, which subsequently form the basis for rule generation and deployment.

C. Risk-Aware Policy Deployment

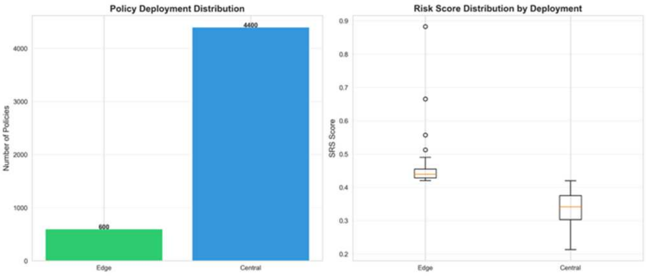

Fig. 5. Policy Deployment Distribution and SRS Score Distribution

The policy deployment simulation examined how extracted tokens and their associated SRS values translate into concrete enforcement rules. Fig. 4 shows that approximately about 12% were deployed to the edge layer, representing the highest-risk subset of rules. The remaining 4,400 policies (about 88%) were allocated to the central enforcement layer.

The distribution of SRS scores between the two layers further reinforces this separation. Edge policies exhibit noticeably higher median SRS values, along with a broader upper tail reflecting their increased likelihood of indicating harmful activity. Central policies, by contrast, cluster around lower SRS values and display a narrower distribution, consistent with lower operational urgency.

This stratified distribution demonstrates that the system successfully identifies high-impact threats and places their corresponding enforcement rules at the most responsive point in the network. By prioritizing high-risk tokens at the edge, the framework reduces detection latency while minimizing unnecessary load on edge computing resources.

D. Summary of HCAPO Performance

TABLE I. OVERALL PERFORMANCE OF THE PROPOSED HCAPO

Metrics	Score
Accuracy	0.8530
Precision (Macro)	0.8203
Recall (Macro)	0.8551
F1-Score (Macro)	0.8348
Precision (Weighted)	0.8620
Recall (Weighted)	0.8530
F1-Score (Weighted)	0.8552

Table I presents the overall performance of the proposed HCAPO model across standard multi-class evaluation metrics. The model achieves an accuracy of 0.853, indicating that more than 85% of all DNS flows are correctly classified across benign, malware, phishing, and spam categories. Macro-averaged precision, recall, and F1-score fall within the 0.82–0.86 range, demonstrating consistent performance even under class imbalance. Macro precision of 0.8203 reflects the model's balanced ability to avoid false positives across different malicious categories, while the macro recall of 0.8551 suggests a strong capability to detect true malicious events irrespective of class frequency. The macro F1-score of 0.8348 further confirms that this balance between precision and recall is maintained for all classes. And weighted metrics provide an additional perspective by incorporating class proportions into the evaluation. The weighted precision of 0.8620 and weighted recall of 0.8530 show that the model performs especially well on high-frequency classes without sacrificing minority-class detection. The weighted F1-score of 0.8552 aligns with this trend, illustrating the stability of the model.

VI. Conclusion

This study introduced HCAPO, a hybrid-cloud security framework that integrates sequential DNS-flow modeling, token-level risk analysis, and risk-aware policy deployment. By leveraging a transformer-based architecture that can capture the temporal and structural characteristics of DNS behavior, the system effectively distinguishes benign traffic from multiple categories of malicious activity. The model achieves an overall accuracy of 0.8530 and maintains balanced recall across malware, phishing, and spam classes, indicating robust performance even under class imbalance. Furthermore, the token extraction mechanism provides a fine-grained representation of threat-inducing domain components, enabling automated generation of security rules that reflect the underlying semantic structure of malicious DNS patterns. The evaluation of the policy deployment component demonstrates that HCAPO can effectively differentiate between high-impact indicators and low-risk signals and allocate enforcement rules accordingly. High-SRS tokens were consistently placed at the edge layer, where rapid mitigation is essential, while the central layer absorbed lower-risk patterns that benefit from deeper analysis. This hierarchical deployment strategy highlights the potential of risk-aware policy placement to improve threat responsiveness while minimizing unnecessary computational overhead.

VII. Future Work

While the results validate the feasibility and effectiveness of the proposed approach, several avenues for further advancement remain. First, the current evaluation focuses on DNS-based telemetry; extending the framework to encompass multi-protocol traffic—such as HTTP, TLS, and email metadata—would provide a more holistic view of modern attack surfaces. Second, a real-world deployment study within enterprise-scale hybrid-cloud infrastructures would enable assessment of operational factors such as policy propagation latency, system scalability, and interactions with existing security appliances. Finally, future work may explore joint optimization techniques that directly couple model predictions, token extraction quality, and deployment strategy, enabling end-to-end learning of security policies informed by real-time environmental feedback.

Acknowledgment

This work was supported by the Institute of Information & Communications Technology Planning & Evaluation(IITP)-Innovative Human Resource Development for Local Intellectualization program grant funded by the Korea government(MSIT)(IITP-2025-RS-2022-00156287) (34%). This research was supported by Basic Science Research Program through the National Research Foundation of Korea(NRF) funded by the Ministry of Education(RS-2025-25398164)(33%). This work was supported by the Technology development Program(RS-2025-25443207) funded by the Ministry of SMEs and Startups(MSS, Korea)(33%)

References

[1] D. Sitaram et al., "Orchestration Based Hybrid or Multi Clouds and Interoperability Standardization," 2018 IEEE International Conference on Cloud Computing in Emerging Markets (CCEM), Bangalore, India, 2018, pp. 67-71

[2] Y. Mansouri, V. Prokhorenko, and M. A. Babar, "An Automated Implementation of Hybrid Cloud for Performance Evaluation of Distributed Databases," J.Netw. Comput. Appl., vol. 167, p. 102740, Oct. 2020

[3] KZ. Li, Q. Lu, L. Zhu, X. Xu, Y. Liu and W. Zhang, "An Empirical Study of Cloud API Issues," in IEEE Cloud Computing, vol. 5, no. 2, pp. 58-72.

[4] LEE, Yungee, et al. Assessing the impact of dos attacks on iot gateway. In: International Conference on Multimedia and Ubiquitous Engineering. Singapore: Springer Singapore, 2017. p. 252-257.

[5] S. B. Mallisetty, G. A. Tripuramallu, K. Kamada, P. Devineni, S. Kavitha and A. V. P. Krishna, "A Review on Cloud Security and Its Challenges," 2023 International Conference on Intelligent Data Communication Technologies and Internet of Things (IDCIoT), Bengaluru, India, 2023, pp. 798-804.

[6] G. Raktate, K. Shelar, P. Parjane, S. Pangavhane, S. More and S. R. Deshmukh, "A Survey on Security Issues and Challenges in Cloud Computing," 2024 International Conference on Decision Aid Sciences and Applications (DASA), Manama, Bahrain, 2024, pp. 1-5.

[7] I A. Mishra, P. Sarat and R. Afza, "A factual study on hybrid multi cloud cyber security threats and proposed methodologies to enable cyber resilience," 2024 IEEE International Conference on Electronics, Computing and Communication Technologies (CONECCT), Bangalore, India, 2024, pp. 1-6.

[8] Sungwoong Yeom, Chulwoong Choi, and Kyungbaek Kim. 2021. AutoEncoder Based Feature Extraction for Multi-Malicious Traffic Classification. In The 9th International Conference on Smart Media and Applications (SMA 2020). Association for Computing Machinery, New York, NY, USA, 285–287.

[9] Yeom, Sungwoong, and Kyungbaek Kim. "Detail analysis on machine learning based malicious network traffic classification." Proc. Int. Conf. Smart Media Appl. 2019.

[10] C. Anjani, R. M. Balajee, G. Divya, Y. S. Sree, K. Padmanabham and S. S. Srithar, "Evolving Threats and AI Solutions for Modern Hybrid Cloud Architectures," 2023 2nd International Conference on Automation, Computing and Renewable Systems (ICACRS), Pudukkottai, India, 2023, pp. 478-484.

[11] S. Yeom and K. Kim, "Improving Performance of Collaborative Source-Side DDoS Attack Detection," 2020 21st Asia-Pacific Network Operations and Management Symposium (APNOMS), Daegu, Korea (South), 2020, pp. 239-242.

[12] D. K. Seth, K. K. Ratra and A. P. Sundareswaran, "AI and Generative AI-Driven Automation for Multi-Cloud and Hybrid Cloud Architectures: Enhancing Security, Performance, and Operational Efficiency," 2025 IEEE 15th Annual Computing and Communication Workshop and Conference (CCWC), Las Vegas, NV, USA, 2025, pp. 00784-00793.

[13] A. Paul, R. Manoj and U. S, "Amazon Web Services Cloud Compliance Automation with Open Policy Agent," 2024 International Conference on Expert Clouds and Applications (ICOECA), Bengaluru, India, 2024, pp. 313-317.

[14] G. Kim, Y. Kim, E. Lee, H. Jang and K. Kim, "Edge-Based Policy Caching for Low Latency Security Enforcement in Hybrid Clouds," 2025 25th Asia-Pacific Network Operations and Management Symposium (APNOMS), Kaohsiung, Taiwan, 2025, pp. 1-6.

[15] Nguyen, Sinh-Ngoc, et al. "Source-side detection of drdos attack request with traffic-aware adaptive threshold." IEICE TRANSACTIONS on Information and Systems 101.6 (2018): 1686-1690.

[16] Max Landauer, Sebastian Onder, Florian Skopik, Markus Wurzenberger, Deep learning for anomaly detection in log data: A survey, Machine Learning with Applications, Volume 12, 2023, 100470, ISSN 2666-8270, https://doi.org/10.1016/j.mlwa.2023.100470.

[17] Guo, H., Lin, X., Yang, J., Zhuang, Y., Bai, J., Zheng, T., ... & Li, Z. (2021). Translog: A unified transformer-based framework for log anomaly detection. arXiv preprint arXiv:2201.00016.

[18] Pospieszny, P., Mormul, W., Szyndler, K., & Kumar, S. (2025). ADALog: Adaptive Unsupervised Anomaly detection in Logs with Self-attention Masked Language Model. arXiv preprint arXiv:2505.13496.