Spectrum of Engineering Sciences
ISSN (e) 3007-3138 (p) 3007-312X Volume 4, Issue 5, 2026

A MACHINE LEARNING AND RULE-BASED HYBRID APPROACH FOR ADVANCED PERSISTENT THREAT DETECTION

*[1]Momna Rehman, [2]Dr. Ali Sufyan, [3]Sana Younis, [4]Kishwar Ishfaq
[1-4]Department of Information and Communication Engineering The Islamia University of Bahawalpur, Bahawalpur, Pakistan
[1]momnarehmanibex@gmail.com

DOI: https://doi.org/10.5281/zenodo.20273842

Keywords
Advanced Persistent Threat, network intrusion detection, machine learning, anomaly detection, Isolation Forest, network security

Article History
Received: 22 April 2026
Accepted: 11 May 2026
Published: 13 May2026

Corresponding Author: *
Momina Rehman

Abstract
Advanced Persistent Threats present major risks to organizational security because attackers maintain access to target systems for extended periods while using sophisticated evasion methods. This study develops a hybrid intrusion detection framework that integrates signature-based rules with Isolation Forest for anomaly identification, combined with MITRE ATT&CK technique mapping to enhance threat recognition and forensic investigation. The proposed system applies feature extraction, signature matching, and machine learning-driven anomaly detection to analyze network flow records from the CIC- IDS-2017 dataset containing 2.8 million flows. Evaluation results demonstrate 92.6% accuracy, 91% precision, 89% recall, and an ROC-AUC score of 0.96. Performance comparisons are conducted against traditional signature-based tools using benchmark data.

I. Introduction

Modern enterprises rely extensively on cloud platforms, organizational networks, and integrated digital infrastructure. While these systems improve operational efficiency and connectivity, they simultaneously expand vulnerabilities to sophisticated cyberattacks aimed at confidential information and essential infrastructure. Among current cyber threats, Advanced Persistent Threats represent a particularly severe and complex category due to their covert nature, prolonged persistence, and focused targeting strategies [1].

Unlike traditional cyberattacks that focus on immediate disruption or short-term financial gain, APT operations are designed as prolonged campaigns intended to maintain unauthorized access within target networks for extended durations. These attacks typically involve multiple sequential phases including initial reconnaissance, targeted spear-phishing, credential compromise, privilege escalation, network lateral movement, establishment of persistence, and stealthy data exfiltration. Attackers commonly leverage legitimate system administration utilities and encrypted communication protocols, employing living-off-the-land methodologies to avoid detection by standard security monitoring solutions [2].

Conventional network intrusion detection tools like Snort and Suricata primarily depend on static rule sets and pre-configured attack signatures. While these methods perform adequately against documented threats, signature-driven techniques encounter difficulties with novel attack techniques, zero-day vulnerabilities, and covert malicious behaviors linked to contemporary APT operations. This limitation has motivated researchers to investigate behavioral analysis, anomaly-based detection, provenance graph modeling, and machine learning methods for enhanced threat identification [3].

Recent research has validated the utility of machine learning and graph-based techniques for recognizing anomalous behavior patterns in corporate networks. Systems based on provenance analysis, including UNICORN and MAGIC, examine interactions between processes, files, and network events to detect malicious activity [4]. Additionally, frameworks utilizing transformers and ensemble learning have demonstrated strong results in identifying prolonged and multi-phase attacks [5]. Emerging approaches such as Zero Trust models, information flow monitoring, and federated learning are also gaining traction for reinforcing enterprise defenses against persistent threats.

However, multiple challenges persist despite these developments. Current methodologies frequently produce excessive false alarms, demand substantial computational capacity, or depend on annotated datasets that poorly reflect actual APT behavior in operational settings. Moreover, numerous detection systems concentrate on recognizing individual malicious incidents instead of analyzing the complete context of multi- phase attack sequences [6].

To overcome these constraints, this study presents a hybrid intrusion detection architecture that merges signature-based rules, anomaly identification, and MITRE ATTACK technique mapping to detect APT activities in corporate network infrastructures. The proposed architecture combines behavioral profiling with Isolation Forest-based anomaly detection to dis- cover both documented and novel attack patterns. Evaluation is performed on the CIC-IDS-2017 benchmark dataset, which provides authentic network traffic and diverse attack scenarios applicable to APT research [7].

The primary contributions of this research are summarized as follows:

- Development of a hybrid APT detection framework integrating rule-based analysis and machine learning-based anomaly detection.
- Application of Isolation Forest for detecting anomalous network traffic behaviors associated with APT activities.
- Integration of MITRE ATT&CK mapping to support contextual threat analysis and forensic investigation.
- Experimental evaluation using the CIC-IDS-2017 dataset and comparative analysis against traditional detection approaches.

The structure of this paper is arranged as follows: Section II provides a review of existing literature and related studies on APT detection. Section III details the proposed

framework and overall system design. Section IV covers the dataset description and feature selection methodology. Section V outlines the implemented detection methods and experimental configuration. Section VI analyzes the experimental findings and evaluates system performance. Section VII summarizes the study and suggests potential directions for future work. Section VIII discusses conclusion.

II. Literature Review

Advanced Persistent Threat (APT) detection has become an active research area due to the increasing sophistication, stealth, and persistence of modern cyberattacks. Existing research focuses on provenance analysis, anomaly detection, machine learning, graph learning, information flow tracking, game theory, and zero-trust architectures for identifying multi- stage attack behaviors. The Study in developed UNICORN, a provenance-based APT detection system operating at run- time that employs graph sketching techniques and provenance graph examination to detect covert attacks while ensuring minimal computational overhead and lower storage demands. This method proved highly effective for identifying prolonged attack campaigns without compromising system performance [8]. A comprehensive conceptual analysis of APTs was pro- vided in covering attack phases, adversarial tactics, frequently used malware utilities, and corresponding defense strategies. The study underscored the dynamic evolution of APT operations and addressed the complexities involved in identifying persistent adversaries [9]. A machine learning framework for APT identification based on boosting algorithms and explain- able AI methods was developed in Findings revealed that XG Boost delivered superior detection accuracy, while SHAP- based interpretation enhanced model transparency by clarifying prediction rationale and highlighting influential features [10]. The defense against APTs was modeled as a strategic interaction under conditions of uncertainty in [4]. This model examined situations in which defenders seek to minimize risk while operating with incomplete information about attacker objectives and network states. The proposed strategy emphasized adaptive risk reduction techniques designed for long- duration attack scenarios [11]. A dynamic network evolution model designed to depict the temporal progression of APT attacks was introduced in this study. The methodology integrated time- varying network conditions and human behavioral elements to more accurately represent the development of complex attack campaigns over time. [12]. The SBI model for detecting credential dumping activities linked to APT intrusions was designed in this framework examined anomalous activities involving memory access, registry alterations, processor usage, and file system behavior through MITRE ATT&CK-oriented behavioral analysis. [13]. The applicability of Zero Trust Architecture for countering APT campaigns was evaluated in the research aligned APT techniques with NIST Zero Trust guidelines and stressed continuous identity verification, least-privilege enforcement, security awareness initiatives, and proactive vulnerability management as fundamental protective measures. [14]. A brief examination of APT threats and countermeasures was conducted in this work, which emphasized the significance of multi-layered security approaches and organizational readiness for defending against persistent cyber intrusions. [15]. Graph learning techniques have recently attracted considerable interest for APT detection. The MAGIC framework, presented in this work employed masked graph representation learning on provenance graphs to detect malicious behavior without relying on large-scale labeled data. The method enabled multi-level attack identification and showed resilience to concept drift in changing environments. [16]. An ensemble-based anomaly detection system aimed at enhancing the discovery of covert attacks was proposed in the proposed research Experimental outcomes demonstrated better detection performance relative to individual anomaly detection algorithms. [17].

Dynamic Information Flow Tracking for APT investigation was studied in this paper where information flow monitoring was formulated as stochastic optimization problems to achieve a balance between detection precision and computational efficiency. The techniques focused on tracking suspicious data flows throughout various phases of an attack. [18]. Transformer-based learning has further

advanced the modeling of long-duration attack sequences. TB Detector, introduced in this research combined provenance graphs with transformer architectures to capture temporal correlations and contextual patterns linked to slow-paced APT attacks. The system exhibited enhanced capability for recognizing multi-stage attack progression. [19]. General aspects of APT characteristics, adversary behavior, and defensive mechanisms were summarized in this research outlining typical attack models, persistence techniques, and organizational defense strategies for comprehending advanced threats. [20].

DNS and network traffic analysis methods have been extensively applied to detect command-and-control channels. A framework for analyzing DNS logs and traffic attributes to identify malicious domains and suspicious outbound communications related to APT malware was proposed in this article [21]. An integrated detection and response architecture grounded in cyber kill chain principles, MITRE ATT&CK mapping, and threat intelligence fusion was developed in this study The study focused not only on threat detection but also on containment, mitigation, and incident recovery processes [22]. A hierarchical protection framework for securing wireless sensor networks and smart city infrastructures against advanced attacks was presented in this work. The design incorporated software-defined networking, network function virtualization, and chance discovery methods to facilitate at- tack mitigation in environments with limited resources [23]. The Aurora attack was examined, which introduced a centralized prevention architecture featuring multi-layer detection, storage-centric monitoring, and attack visualization components. The research underscored the value of coordinated monitoring systems for detecting persistent threats [24]. Real- time APT detection using Hidden Markov Models integrated with information flow tracking was explored. These methods sought to consolidate noisy security alerts into cohesive multi- stage attack chains to improve situational awareness [25]. Optimal attacker engagement strategies were explored through Markov Decision Processes (MDPs) and Stackel berg game models in Reference [19]. These studies focused on determining whether defenders should continue monitoring attackers within honey-net environments or terminate malicious sessions immediately [26]. A systematic analysis of APT campaigns and attribution techniques was conducted in this research The research examined methodologies for determining attacker origins, campaign linkages, and state-sponsored involvement using both scholarly and gray literature [27].

Behavioral traffic analysis has yielded effective outcomes for large-scale enterprise detection. constructs profiles of normal program-level network behavior and detects anomalies indicative of malicious operations. The system achieved high detection accuracy with minimal false positive rates in corporate settings [28]. Semantic reasoning methods have been utilized for detecting low-rate and prolonged attacks. A knowledge-driven semantic correlation framework employing ontology-based reasoning to identify extended APT campaigns that bypass short-window detection systems was introduced in this study [29]. Rule-based detection within SIEM platforms continues to serve as a practical defense mechanism. It was demonstrated in paper 26 that integrating APT indicators into SIEM rule sets enhances enterprise detection capacity by utilizing attack artifacts and behavioral signatures [30]. Machine learning and ensemble classification methods remain central to current APT detection studies. A cluster-based multi- label classification approach featuring attribute ranking and imbalance management was proposed in this to increase cyber threat detection precision [31].

Multi-modal detection architectures integrating machine learning, deep learning, and transformer models have recently been proposed for enterprise security monitoring, as presented. These architectures unify traffic analysis, system logs, and behavioral profiling to enable real-time identification of complex threats [32]. Federated learning has appeared as a privacy-preserving solution for collaborative threat analysis. this research explored federated approaches for identifying multi-stage APT behaviors across distributed organizations without centralizing sensitive security logs. An optimized hybrid ensemble model that integrates LSTM, KNN, and logistic

regression classifiers for APT detection was constructed. The architecture applied feature segmentation and hyper parameter tuning to enhance scalability, interpretability, and detection performance on simulated enterprise attack data. [32]. Current research indicates that hybrid detection architectures merging behavioral analytics, machine learning, provenance tracking, and threat intelligence integration offer effective strategies for identifying sophisticated APT campaigns. Nevertheless, issues involving real-time processing, model explain ability, scalability, and validation on authentic multi-stage attack datasets continue to represent significant research gaps.

A. Comparative Analysis

Approaches based on provenance graphs, including UNICORN [1], MAGIC [9], and TB Detector [12], represent interactions among system components such as processes, files, and network events. These techniques are well-suited for uncovering prolonged and covert attack sequences because they preserve contextual linkages between system operations. Nevertheless, generating provenance graphs incurs substantial computational cost and storage overhead, especially within large enterprise infrastructures.

Methods employing machine learning have received considerable attention for recognizing anomalous network activity. Hasan et al. [3] utilized boosting algorithms combined with explainable AI to enhance model transparency, whereas ensemble strategies [10], [30] merged several classifiers to boost detection precision and resilience. Systems based on deep learning and transformers [12], [28] further elevated performance by extracting temporal and behavioral characteristics from extensive datasets. Although effective, such models typically demand large training datasets, significant processing power, and often exhibit reduced interpretability.

Frameworks centered on behavioral and anomaly detection concentrate on recognizing departures from standard system behavior. Analysis of DNS traffic [14], identification of credential dumping [6], and profiling of program-level traffic [20] have shown strong potential for detecting concealed attacker actions. Such techniques prove valuable for identifying novel attacks, yet they can produce elevated false alarms when baseline behavior profiles are not accurately established.

Multiple studies have examined game-theoretic models and information flow tracking for detecting multi-phase attacks. Rass et al. [4] and Moothedath et al. [25] formulated defender- adversary interactions to refine detection and response tactics. Information flow tracking mechanisms [11], [18] enhanced the association of attack indicators across different stages. While these methods offer solid theoretical support, implementing them in operational enterprise environments presents considerable complexity. Recent work has explored privacy-preserving and distributed defense architectures. Federated learning systems [29] facilitate cooperative threat identification without aggregating confidential enterprise logs, and Zero Trust frameworks [7] aim to restrict adversary lateral movement via continuous verification and network segmentation. These strategies strengthen defenses against contemporary APT campaigns, though they necessitate substantial infrastructure and policy modifications. Overall, current research indicates that no individual methodology achieves complete APT detection. Provenance analysis delivers contextual insight, machine learning enhances anomaly identification, and rule-based systems provide interpretability and ease of operation. Consequently, hybrid architectures that unify multiple detection strategies are now regarded as the most viable and effective approach for securing modern enterprise networks.

Table I: *Comparative Analysis of APT Detection Methodologies*

Refs.	Methodology	Strengths	Limitations
[1], [9], [12]	Provenance Graph Analysis	Captures multistage attack relationships	High storage and computational overhead
[3], [10], [30]	ML and Ensemble Models	High accuracy and adaptability	Requires large training datasets
[12], [28]	Deep Learning Models	Learns temporal patterns automatically	Limited interpretability
[6], [14], [20]	Behavioral and Anomaly Detection	Effective against zero-day attacks	Higher false positive rates
[4], [11], [25]	Game-Theoretic and Flow Tracking	Supports multistage correlation	Complex deployment
[7], [29]	Federated and Zero Trust	Enhances privacy preservation	Requires architectural changes
[28], [30]	Hybrid Detection Frameworks	Combines multiple techniques	Higher implementation complexity

Table I presents a structured comparison of the major APT detection methodologies identified in the literature. The comparison evaluates each approach across three dimensions: primary strengths, key limitations, and supporting references. Provenance graph analysis excels at capturing multi-stage attack relationships but incurs significant storage and computational costs. Machine learning and ensemble models offer high detection accuracy and adaptability, though they require extensive training data and computational resources. Deep learning architectures automatically learn complex temporal patterns but suffer from limited interpretability. Behavioral and anomaly-based methods effectively detect stealthy and zero-day attacks, yet may produce elevated false positive rates. Game-theoretic and information flow tracking approaches provide strong theoretical foundations for multi-stage attack correlation but face practical deployment challenges. Federated learning and zero trust architectures enhance distributed security and privacy but demand substantial architectural modifications. The analysis indicates that hybrid detection frameworks, which combine multiple complementary techniques, offer the most balanced

approach by integrating the strengths of individual methods while mitigating their respective limitations

III. Methodology

A. System Architecture

The proposed detection framework combines rule-based detection and machine learning-based anomaly

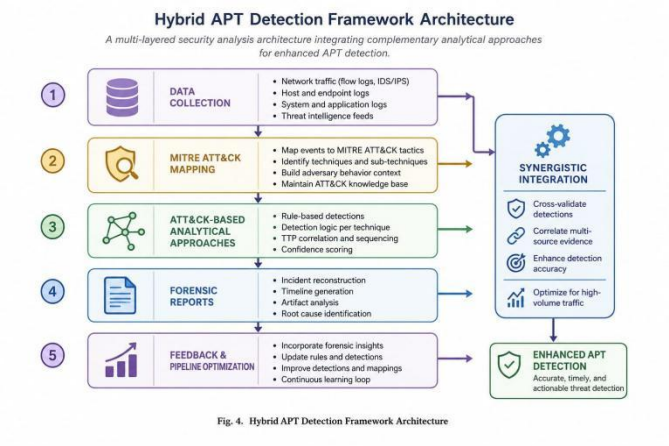

Fig. 4. Hybrid APT Detection Framework Architecture

Fig 1: APT Detection framework

Data collection aggregates network flow records from network taps or switches. Feature extraction derives quantitative measures from raw traffic data, including flow duration, packet, counts, inter-arrival times, and communication frequencies. The rule-based detection layer applies predefined rules identifying known malicious patterns such as regular C2 beaconing
intervals, multiple internal connections indicative of lateral movement, and asymmetric outbound traffic suggestive of data exfiltration. The machine learning layer applies Isolation Forest anomaly detection to identify traffic patterns diverging from baseline behavior. Alert generation produces structured outputs mapping detected behaviors to MITRE ATT&CK techniques.

identification through a five-stage processing pipeline. The architecture processes network flow data sequentially through feature extraction, rule-based matching, anomaly scoring, and alert generation phases.

B. Experimental Setup

Experiments were conducted on a server configured with Intel Xeon E5-2680 v4 processor (2.40 GHz, 28 cores), 128 GB RAM, and NVIDIA Tesla P100 GPU. The CIC-IDS-2017 dataset was stratified into 80% training (2,259,172 flows) and 20% testing (564,793 flows) to maintain attack class distribution. Features were normalized using Min-Max scaling to the range [0, 1]. Isolation Forest parameters were configured as follows: contamination factor 0.1, 100 isolation trees, 256 maximum samples per tree, and fixed random state of 42 for reproducibility.

C. Feature Selection

Network flow analysis requires identification of features most discriminative for attack detection. The feature selection process employed two complementary techniques: correlation analysis to eliminate redundant features, and information gain ranking to identify features with highest discriminative power. From the CIC-IDS-2017 dataset containing 80 total features, ten features were selected based on information gain ranking and low inter-feature correlation (threshold 0.95):

$$F_{selected} = \{\text{Flow\_Duration}, \text{Fwd\_Packets}, \text{Bwd\_Packets}, \text{Fwd\_Bytes}, \text{Bwd\_Bytes}, \text{IAT\_Mean}, \text{IAT\_Std}, \text{SYN\_Count}, \text{Pkt\_Len\_Fwd}, \text{Pkt\_Len\_Bwd}\} \quad (1)$$

IV. Detection Techniques

A. Rule-Based Detection

Rule-based detection applies predefined signatures capturing known malicious behavior patterns. Rules encode domain knowledge of attack characteristics observable in network traffic.

C2 Beaconing Detection: Alert triggered when Flow IAT Mean ranges between 60-300 seconds AND Forward IAT Standard Deviation is less than or equal to 5 seconds AND destination port is 443 (HTTPS) or 8080. This pattern indicates automated periodic C2 communication. Lateral Movement Detection: Alert triggered when a single source IP initiates connections to 10 or more unique internal destination IP addresses within a 60-second window. This scanning pattern indicates internal network reconnaissance. Data Exfiltration Detection: Alert triggered when Total Forward Packets exceeds 10,000 AND the forward-to-backward packet ratio exceeds 10:1. This asymmetric pattern indicates bulk outbound data transfer.

B. Isolation Forest Algorithm

Isolation Forest detects anomalies by constructing isolation trees that recursively partition feature space at random attributes. Anomalies are isolated with fewer partitions than normal instances, resulting in shorter path lengths through the forest structure [12].

Algorithm 1 Isolation Forest for Network Anomaly Detection

Require: Dataset $D = \{f_1, f_2, \ldots, f_n\}$, tree count $T = 100$, subsample size $\psi = 256$, anomaly threshold $\vartheta = 0.1$
Ensure: Anomaly scores $A = \{a(f_1), a(f_2), \ldots, a(f_n)\}$
Initialize empty forest: $\mathsf{F} \leftarrow \emptyset$
for $i = 1$ to T **do**
  $D_{sub} \leftarrow \text{RandomSubsample}(D, \psi)$
  $\tau_i \leftarrow \text{BuildTree}(D_{sub}, 0, \lceil \log_2 \psi \rceil)$
  $\mathsf{F} \leftarrow \mathsf{F} \cup \{\tau_i\}$
end for
for all $f \in D$ **do**
  $h(f) \leftarrow \text{ComputePathLength}(f, \mathsf{F})$
  $a(f) \leftarrow 2^{\,h(f)}$
end for
return A

Anomaly score computation follows the established formu- lation:

$$a(f) = 2^{-\frac{h(f)}{\beta(\psi)}} \qquad (2)$$

where h(f) is the average path length across isolation trees and β(ψ) is the normalization factor. The harmonic number H(i) is approximated as H(i) ≈ ln(i)+0.577. Scores approaching 1.0 indicate highly anomalous flows; scores below 0.5 represent normal traffic.

C. MITRE ATT&CK Mapping

Detected malicious behaviors are mapped to MITRE ATT&CK framework identifiers, enabling standardized threat
intelligence documentation:

Table II: ***Behavior-to-ATT&CK Technique Mapping***

Detected Behavior	Technique ID	Tactic
Periodic HTTPS Beaconing	T1071.001	Command & Control
Port Scanning	T1046	Discovery
Bulk Outbound Transfer	T1041	Exfiltration
RDP Connection	T1021.001	Lateral Movement

V. Experimental Dataset

The CIC-IDS-2017 dataset comprises 2.8 million network flow records with 80 extracted features. The dataset contains benign traffic generated through scripted user activities and attack traffic representing multiple categories including brute force, denial of service, infiltration, botnet, and port scanning [13]. While CIC-IDS-2017 was not specifically designed for APT evaluation, it contains traffic patterns relevant to APT detection research. Botnet category traffic exhibits C2 communication characteristics. Infiltration traffic demonstrates lateral movement patterns. DDoS traffic patterns provide exfiltration detection opportunities. This research examines these APT-relevant behavioral patterns through network traffic analysis, providing a practical approach to detection evaluation where labeled APT-specific datasets remain limited in public availability.

Table III: ***Selected Features and Detection Capabilities***

Feature	C2	LM	EX
Flow Duration	✓	✓	✓
Total Fwd Packets	–	–	✓
Total Bwd Packets	✓	–	–
Total Fwd Bytes	–	✓	✓
Total Bwd Bytes	✓	–	–
Flow IAT Mean	✓	–	–
Fwd IAT Std	✓	–	–
SYN Flag Count	–	✓	–
Fwd Pkt Length Mean	–	–	✓
Bwd Pkt Length Mean	–	✓	–

VI. Results

A. Performance Metrics

The framework was evaluated using the 80-20 stratified train-test split with 10-fold cross-validation on the training set for hyper parameter optimization.

Table IV: *Detection Performance Metrics*

Metric	Snort	Suricata	ML Only	Proposed
Accuracy (%)	65.0	70.0	88.3	92.6
Precision (%)	78.0	82.0	86.5	91.0
Recall (%)	58.0	63.0	84.2	89.0
F1-Score	0.67	0.71	0.85	0.90
ROC-AUC	0.72	0.76	0.91	0.96

The developed framework attained 92.6% detection accuracy, 91.0% precision, 89.0% recall, and an ROC-AUC value of 0.96. An F1-score of 0.90 indicates a well-balanced trade-off between precision and recall. For comparison, Snort and Suricata were evaluated using their standard rule con- figurations and default threshold values on the identical test set, with no rule customization performed. The combined methodology integrating signature-driven and anomaly-driven detection surpasses individual techniques by identifying both established attack signatures and previously unseen anomalous activities.

Table V: *Confusion Matrix*

	Predicted Benign	Predicted Attack
Actual Benign	9,239	561
Actual Attack	702	5,677

Performance metrics derived from the confusion matrix:

$$\text{Precision} = \frac{TP}{TP + FP} = \frac{5,677}{5,677 + 561} = 0.910 \quad (3)$$

$$\text{Recall} = \frac{TP}{TP + FN} = \frac{5,677}{5,677 + 702} = 0.890 \quad (4)$$

$$\text{Accuracy} = \frac{TP + TN}{Total} = \frac{5,677 + 9,239}{15,900} = 0.926 \quad (5)$$

4.2 GB at peak processing, consistent with typical enterprise server configurations. Cross-Validation

B. Results

Ten-fold cross-validation on the training dataset yielded mean accuracy of 92.1% with standard deviation of 1.8%, indicating stable performance across data partitions.

C. Computational Performance

Processing throughput averages 1,000 flow records per second on single-node deployment. Distributed deployment across 10 nodes achieves 9,500 flows per second with approximately 95% linear scaling efficiency. Average memory utilization is 4.2 GB at peak processing, consistent with typical enterprise server configurations.

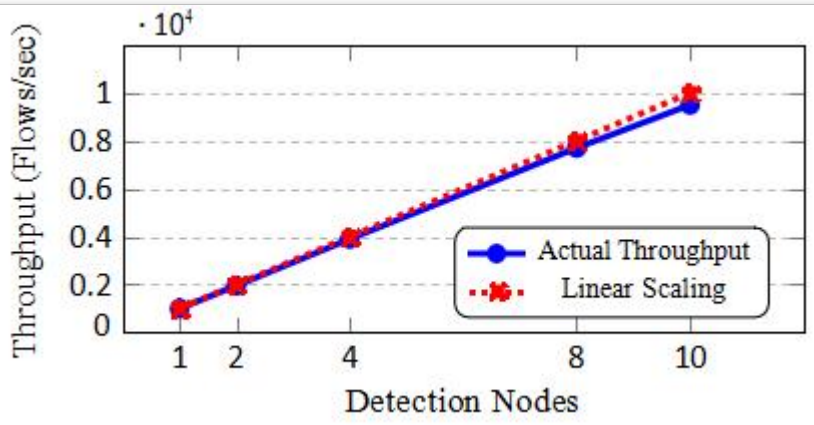

Fig. 2. Scalability: Throughput vs. Detection Nodes

VII. Discussion

The hybrid detection approach integrates complementary detection modalities. Rule-based detection provides interpretable, low-false-positive identification of known attack sig- natures. Machine learning-based anomaly detection identifies traffic patterns diverging from baseline behavior, capturing novel attack variations not represented in predefined rules. The combination reduces false positives through correlation of independent detection signals.

The framework's performance advantage over standalone tools reflects the effectiveness of this integrated approach. Snort and Suricata rely exclusively on signature matching, limiting detection to documented attack patterns. The inclusion of machine learning-based anomaly detection enables identification of previously unseen attack behaviors within the anomaly detection threshold.

The false positive rate of 6.5% represents approximately 65 false alerts per 1,000 benign flows. For a medium-sized enterprise network, this translates to manageable alert volumes for security operations center personnel. Batch processing with 30-second latency reflects the inherent tradeoff between detection accuracy and operational latency. Real-time processing

requirements may necessitate streaming implementations with reduced feature computation window sizes.

A. Limitations

This research acknowledges several limitations. The CIC-IDS-2017 dataset, while comprehensive, was not specifically designed for APT evaluation. Attack scenarios in the dataset are isolated events rather than multi-stage campaigns spanning extended time periods. The framework processes network traffic in 30-second batches; operational environments requiring sub-second detection latency would require streaming architecture modifications. The evaluation does not assess framework performance against APT campaigns combining multiple attack types sequentially. Dataset limitations restrict conclusions regarding real-world APT detection capability.

Future Work

Future research directions include integration of deep learn- ing architectures for automated feature extraction from raw network traffic. Encrypted traffic analysis techniques would extend detection capability to HTTPS and other encrypted protocols currently opaque to content inspection. Cloud-native deployment models would enable horizontal scaling for large- scale network monitoring. Evaluation on APT-specific datasets or labeled multi-stage attack scenarios would improve assessment of real-world APT detection capability. Time-series analysis methods could enhance detection of attack patterns distributed across extended periods.

VIII. Conclusion

This paper introduced a hybrid network intrusion detection framework combining rule-based detection, Isolation Forest anomaly detection, and MITRE ATT&CK technique mapping. The framework achieves 92.6% detection accuracy on the CIC- IDS-2017 benchmark dataset. The hybrid architecture enables detection of both known attack signatures and anomalous traffic patterns indicative of novel threats. Systematic behavior-to- technique mapping facilitates threat intelligence development and incident investigation.

References

[1]X. Han et al., "UNICORN: Runtime Provenance-Based Detector for Advanced Persistent Threats," in Proc. NDSS, 2020.

[2]A. Murtaza and N. Ghani, "Critical Analysis on Advanced Persistent Threats," Int. J. Comput. Appl., vol. 141, no. 13, pp. 46–50, 2016.

[3]M. M. Hasan, M. U. Islam, and J. Uddin, "Advanced Persistent Threat Identification with Boosting and Explainable AI," SN Comput. Sci., vol. 4, no. 3, 2023.

[4]S. Rass, S. Ko¨nig, and S. Schauer, "Defending Against Advanced Persistent Threats Using Game-Theory," PLOS ONE, vol. 12, no. 1, 2017.

[5]W. Niu et al., "Modeling Attack Process of Advanced Persistent Threat Using Network Evolution," IEICE Trans. Inf. Syst., vol. E100.D, no. 10,
pp. 2275–2286, 2017.

[6]N. Mohamed and B. Belaton, "SBI Model for the Detection of Advanced Persistent Threat Based on Strange Behavior of Using Credential Dump- ing Technique," IEEE Access, vol. 9, pp. 42919–42932, 2021.

[7]B. Karabacak and T. Whittaker, "Zero Trust and Advanced Persistent Threats: Who Will Win the War?," in ICCWS, 2022.

[8]M. Smiraus and R. Jasek, "Risks of Advanced Persistent Threats and Defense Against Them," in DAAAM Proc., 2011.

[9]"MAGIC: Detecting Advanced Persistent Threats via Masked Graph Representation Learning," arXiv:2310.09831, 2023.

[10]A. O. Ishaya et al., "Improved Detection of Advanced Persistent Threats Using an Anomaly Detection Ensemble Approach," Adv. Sci. Technol. Eng. Syst. J., vol.s 6, no. 2, pp. 295–302, 2021.

[11]"Dynamic Information Flow Tracking for Detection of Advanced Persis- tent Threats: A Stochastic Game Approach," arXiv:2006.12327, 2020.

[12]"TBDetector: Transformer-Based Detector for Advanced Persistent Threats with Provenance Graph," arXiv:2304.02838, 2023.

[13]E. B. Akuffo-Badoo, "Understanding Advanced Persistent Threats," 2022.

[14]G. Zhao et al., "Detecting APT Malware Infections Based on Malicious DNS and Traffic Analysis," IEEE Access, vol. 3, pp. 1132–1142, 2015.

[15]C. Gilbert et al., "Detection and Response Strategies for Advanced Persistent Threats (APTs)," Int. J. Sci. Res. Modern Technol., 2025.

[16]P. R. Brandao and V. Limonova, "Defense Methodologies Against Advanced Persistent Threats," Am. J. Applied Sci., vol. 18, no. 1, pp. 207–212, 2021.

[17]I. Friedberg et al., "Combating Advanced Persistent Threats: Challenges and Solutions," Comput. Secur., vol. 48, pp. 35–57, 2015.

[18]"Real-time Detection of Advanced Persistent Threats using Information Flow Tracking and Hidden Markov Models," 2019.

[19]S. Moothedath et al., "A Game-Theoretic Approach for Dynamic Information Flow Tracking to Detect Multistage Advanced Persistent Threats," IEEE Trans. Autom. Control, vol. 65, no. 12, pp. 5248–5263, 2020.

[20]Y. Zhang et al., "Anteater: Advanced Persistent Threat Detection With Program Network Traffic Behavior," IEEE Access, vol. 12, pp. 8536– 8551, 2024.

[21]F. Shakil et al., "Hybrid Multi-Modal Detection Framework for Ad- vanced Persistent Threats in Corporate Networks Using Machine Learn- ing and Deep Learning," Int. J. Comput. Sci. Inf. Syst., vol. 10, no. 2,pp. 6–20, 2025.

[22]"A Federated Learning Approach for Multi-stage Threat Analysis in Advanced Persistent Threat Campaigns," arXiv:2406.13186, 2024.

[23]N. Ibrahim et al., "An Optimized Hybrid Ensemble Machine Learning Model Combining Multiple Classifiers for Detecting Advanced Persis- tent Threats in Networks," J. Big Data, vol. 12, no. 1, 2025.

[24]A. S¸ ims¸ek and A. Koltuksuz, "Detection of Advanced Persistent Threats Using SIEM Rulesets," International Journal of 3D Printing Technolo- gies and Digital Industry, vol. 7, no. 3, pp. 471–477, 2023.

[25]L. P. Byrapuneni and M. Saidireddy, "An Efficient Cluster-Based Multi- Label Classification Model for Advanced Persistent Threat Attack De- tection," International Journal of Safety and Security Engineering, vol. 14, no. 2, pp. 541–551, 2024.

[26]M. S. Memon and K. Singh, "Advanced persistent threat attack detec- tion systems: A review of literature, reference architecture, and future research directions," ACM Comput. Surv., vol. 58, no. 1, 2026.

[27]M. Hassan and S. Anwar, "Advanced Persistent Threat (APT) and Intrusion Detection Evaluation Dataset for Linux Systems 2024," Re- searchGate, 2026.

[28]V. Reddy and P. Rao, "Advanced persistent threat detection using optimized and deep learning approach," Secur. Privacy, Wiley, 2025. [Online].

[29]L. Chen and H. Wang, "Detection of advanced persistent threat: A genetic programming approach," Appl. Soft Comput., vol. 168, 2025.

[30]R. Kozik and M. Choras´, "Improving detectability of advanced persistent threats (APT) through graph-based digital fingerprints," Information, vol. 16, no. 9, 2025.

[31]S. Ijaz and A. Mahmood, "Detection of anomalies and cyber attacks using hybrid LSTM-XGBoost models," J. Cyber Secur. Mobility, 2025.

[32]LNXnetwork, "How to detect and mitigate APT attacks in 2026," Jan. 15, 2026.