



# NIDS — Network Intrusion Detection System: Design and Implementation of a Hybrid Signature-Based and Anomaly-Based Real-Time Network Monitoring Platform

S.Akash

III BCA C Student

Department of Computer Applications  
VISTAS, Chennai

Dr. K.Dharmarajan Msc,M phil,Phd

Professor

Department of Computer Applications  
VISTAS, Chennai

**Abstract:** A Network Intrusion Detection System (NIDS) is a critical cybersecurity solution designed to monitor, analyze, and detect unauthorized or malicious activities within a network. With the rapid growth of digital infrastructure and internet-based services, networks are increasingly vulnerable to cyber threats such as malware, denial-of-service attacks, and data breaches. This project focuses on the design and implementation of an efficient NIDS that leverages both signature-based and anomaly-based detection techniques. The system continuously inspects network traffic in real time, identifying suspicious patterns and comparing them against a database of known attack signatures. Additionally, machine learning algorithms can be integrated to detect unknown or emerging threats by analyzing deviations from normal network behavior. The proposed system includes alert mechanisms that notify administrators of potential intrusions, enabling quick response and mitigation. Key features of the NIDS include high detection accuracy, low false-positive rates, scalability, and real-time processing capabilities. Security measures such as encryption and secure logging are incorporated to ensure data integrity and confidentiality. Overall, this project demonstrates the importance of proactive network monitoring and intelligent threat detection in safeguarding digital systems and maintaining secure communication environments.

**Keywords** — Network Intrusion Detection System, Signature-Based Detection, Anomaly-Based Detection, Machine Learning, Real-Time Traffic Analysis, Deep Packet Inspection, Snort, Suricata, Alert Mechanism, Cybersecurity, Encryption, Secure Logging.

## I. INTRODUCTION

The explosive proliferation of networked devices, cloud-based services, and internet-dependent applications has fundamentally transformed the threat landscape facing modern organizations. Every day, enterprise networks process billions of packets carrying sensitive financial data, intellectual property, personal information, and critical infrastructure commands. The same connectivity that enables this operational richness simultaneously creates vast attack surfaces exploitable by adversaries ranging from opportunistic automated bots to sophisticated nation-state threat actors.

Traditional perimeter defenses — firewalls, access control lists, and VPN gateways — establish a barrier at the network boundary but provide limited visibility into traffic flowing within the perimeter. An attacker who

successfully circumvents or is already inside the perimeter can move laterally, escalate privileges, and exfiltrate data for extended periods before detection. Industry reports consistently indicate that the mean time to identify a breach exceeds 200 days in many organizational environments, a figure that underscores the inadequacy of perimeter-only security models.

Network Intrusion Detection Systems (NIDS) address this gap by providing continuous, real-time inspection of network traffic at strategic monitoring points. A NIDS analyses packet headers and payloads, reconstructs protocol sessions, and applies detection logic to identify patterns consistent with known or suspected attack behaviour. When an intrusion is detected, the system generates alerts that notify security administrators, enabling timely investigation and response before significant damage occurs.

This paper presents the design and implementation of a hybrid NIDS that combines signature-based detection — pattern matching against a curated database of known attack signatures — with anomaly-based detection powered by machine learning. The hybrid approach overcomes the complementary weaknesses of each method: signature-based detection achieves high precision on known threats but cannot detect novel attacks, while anomaly-based detection identifies deviations from normal behaviour that may indicate zero-day exploits but tends toward higher false-positive rates. By combining both methods within a unified detection engine, the proposed system achieves superior coverage across the full spectrum of network threats.

The remainder of this paper is organized as follows: Section II reviews related work in intrusion detection research. Section III describes the limitations of existing systems. Section IV presents the system architecture. Section V details the proposed methodology. Section VI describes the functional modules. Section VII covers implementation details. Section VIII presents testing and evaluation results. Section IX concludes the paper.

## II. RELATED WORK

Intrusion detection has been an active research domain since the pioneering work of Anderson (1980), who first proposed the concept of auditing computer usage to detect anomalous behaviour indicative of security breaches. Denning (1987) formalized the statistical anomaly detection model, establishing the theoretical foundation for comparing current system behaviour against a learned baseline and flagging deviations above a threshold. These early models, while conceptually sound, were limited by the computational constraints of their era and the absence of large labelled datasets for training and evaluation.

The introduction of the KDD Cup 1999 dataset provided a standardized benchmark for evaluating network intrusion detection algorithms, stimulating extensive comparative research. Tavallaee et al. (2009) identified significant statistical deficiencies in the KDD Cup dataset and introduced the NSL-KDD dataset as a corrected alternative, which has since become the most widely used benchmark in the field. The CICIDS 2017 and 2018 datasets, developed by the Canadian Institute for Cybersecurity, provide more realistic modern traffic captures including contemporary attack categories such as infiltration, botnet activity, and web-based attacks.

Signature-based NIDS implementations have been extensively studied through the Snort and Suricata open-source platforms. Roesch (1999) introduced Snort as a lightweight packet sniffer and intrusion detector capable of real-time traffic analysis against a rule-based signature library. Subsequent work by Caswell et al. demonstrated that the Snort rule language provides expressive power sufficient to describe complex multi-packet attack patterns while maintaining throughput compatible with gigabit network links. Suricata, introduced by the Open Information Security Foundation, extended the Snort model with native multi-threaded processing, enabling full line-rate inspection on multi-core hardware.

Machine learning approaches to anomaly detection have proliferated following the availability of large-scale network traffic datasets and accessible ML frameworks. Buczak and Guven (2016) conducted a comprehensive survey of machine learning methods applied to intrusion detection, evaluating decision trees, support vector machines, naive Bayes classifiers, and artificial neural networks across multiple datasets. Their analysis demonstrated that ensemble methods — particularly Random Forest — consistently achieve the highest detection rates with acceptable false-positive rates across diverse attack categories. Subsequent deep learning approaches, including LSTM-based sequence models for temporal traffic pattern analysis and autoencoder-based anomaly detectors, have demonstrated further improvements in detecting sophisticated multi-stage attacks.

The integration of signature-based and anomaly-based detection within a unified hybrid framework has been explored by several researchers. Mukherjee et al. proposed a layered architecture in which signature matching

provides a fast first-pass filter that diverts high-confidence known-attack traffic to alert generation while forwarding uncertain traffic to a slower but more expressive anomaly detection engine. This architecture achieves the latency characteristics of signature-based systems for the majority of traffic while preserving anomaly detection coverage for novel threats — the same principle underlying the proposed system.

### III. EXISTING SYSTEM

Existing network intrusion detection deployments fall into three broad categories, each exhibiting characteristic limitations that motivate the hybrid approach proposed in this paper.

Pure signature-based systems, exemplified by production Snort and Suricata deployments, rely entirely on manually crafted rule sets maintained by security vendors and community contributors. While these systems achieve extremely low false-positive rates on known attack patterns — because rules are authored with precise specificity — they are fundamentally reactive: a new attack variant cannot be detected until a corresponding signature is written, reviewed, distributed, and loaded by the deployed sensor. The mean time from the identification of a novel attack to the availability of a corresponding Snort rule has historically ranged from hours to days, during which deployments using only signature-based detection remain blind to the new threat. Furthermore, signature databases must be continuously updated, requiring ongoing operational maintenance and network connectivity to update distribution servers.

Pure anomaly-based systems address the zero-day detection gap but introduce the significant operational burden of elevated false-positive rates. A statistical anomaly detector trained on a baseline of normal traffic will raise alerts for any deviation exceeding its threshold — including legitimate but unusual user behaviour, new application deployments, scheduled maintenance activities, and seasonal traffic pattern shifts. Security operations centres receiving high volumes of false-positive alerts suffer from analyst fatigue, which leads to slower response times and, paradoxically, an increased risk that genuine intrusion alerts are overlooked or dismissed.

Commercially available NIDS appliances from vendors such as Cisco, Palo Alto Networks, and IBM Security QRadar address these limitations through proprietary hybrid architectures, threat intelligence feeds, and automated response capabilities. However, these solutions carry significant licensing costs that place them beyond the reach of small and medium enterprises, academic institutions, and government agencies with constrained security budgets. They also introduce vendor dependency and limit the transparency and customizability of detection logic. The limitations of existing systems can be summarized as follows:

- Signature-only systems cannot detect zero-day exploits or novel attack variants
- Anomaly-only systems generate excessive false-positive alerts that overwhelm security analysts
- Commercial hybrid solutions are cost-prohibitive for budget-constrained organizations
- Proprietary systems lack transparency in detection logic, limiting auditability and customization
- Most existing systems lack integrated encryption and secure logging for compliance requirements
- Few open-source solutions provide unified management dashboards with real-time alert visualization

### IV. SYSTEM ARCHITECTURE

The proposed NIDS follows a five-layer pipeline architecture that separates the concerns of traffic capture, preprocessing, detection, alerting, and management into distinct, independently scalable components. The architecture is illustrated conceptually as a sequential processing pipeline through which each captured packet flows from ingestion to final disposition.

#### Traffic Capture Layer

The capture layer is responsible for acquiring raw network packets from one or more network interface cards (NICs) configured in promiscuous mode. The system uses libpcap on Linux and WinPcap/Npcap on Windows to achieve kernel-bypass packet capture, minimising the CPU overhead of copying packet data between kernel and user space. Captured packets are timestamped with nanosecond precision and placed in a lock-free ring buffer for consumption by the preprocessing layer. The capture layer supports both online (live traffic) and offline (PCAP file replay) modes, enabling analysis of historical traffic captures for forensic investigation.

### Preprocessing and Protocol Analysis Layer

The preprocessing layer reassembles TCP streams from individual packets, decodes application-layer protocols, and extracts structured feature vectors for consumption by both detection engines. Protocol decoders are implemented for HTTP, HTTPS (via TLS handshake metadata extraction), DNS, SMTP, FTP, SSH, and generic TCP/UDP flows. For each completed flow or configurable sliding time window, the layer computes a feature vector comprising 41 statistical attributes including packet inter-arrival times, byte rate, protocol distribution, flag counts, and payload entropy. These feature vectors are the input to the anomaly detection engine.

### Signature Detection Engine

The signature engine implements a variant of the Aho-Corasick multi-pattern string matching algorithm, enabling simultaneous matching of thousands of attack signatures against packet payloads in a single linear scan. Signatures are stored in a compiled rule database that encodes protocol context, directional constraints, content matching expressions, and threshold conditions. The engine processes each packet against all applicable rules in  $O(n)$  time relative to packet length, independent of the number of active signatures. Rule updates are applied atomically using a double-buffering scheme that prevents detection gaps during rule set reloads.

### Anomaly Detection Engine

The anomaly detection engine implements a two-stage model: a Random Forest classifier trained offline on the NSL-KDD and CICIDS 2017 datasets provides the primary anomaly score, while an Isolation Forest unsupervised detector provides a complementary score for traffic patterns that diverge from the Random Forest training distribution. The two scores are combined using a weighted linear fusion function, with weights tuned to minimise the combined false-positive and false-negative rates on a held-out validation dataset. The engine is loaded as a serialised scikit-learn pipeline at startup and requires no retraining during normal operation.

### Alert and Response Layer

Detections from both engines are forwarded to the alert layer, which applies deduplication logic to suppress redundant alerts from the same source within a configurable suppression window. Unique alerts are enriched with contextual metadata — geolocation of source IP, associated CVE references for signature matches, and confidence scores for anomaly detections — before being written to the encrypted alert database and forwarded to the management dashboard via a WebSocket push interface. The alert layer also supports configurable automated responses including firewall rule insertion via the iptables API and SNMP trap generation for integration with existing network management systems.

### Management Dashboard

The management dashboard is a web application implemented in Flask and React that provides real-time alert visualization, traffic statistics, rule management, and system health monitoring. Role-based access control restricts rule modification and alert acknowledgement to authorized security administrators. All dashboard communications are secured with TLS 1.3 and authenticated via JWT tokens with configurable expiry periods.

*Figure 1: NIDS Architecture — Traffic Capture → Preprocessing → Signature Engine + Anomaly Engine → Alert Fusion → Dashboard*

## V. PROPOSED METHODOLOGY

The methodology of the proposed NIDS is grounded in a hybrid detection philosophy that exploits the complementary strengths of rule-based pattern matching and data-driven statistical modelling. The methodology proceeds through four phases: data collection and labelling, model training and validation, system integration and deployment, and continuous operational refinement.

In the data collection phase, network traffic captures from the NSL-KDD, CICIDS 2017, and CICIDS 2018 datasets are preprocessed to extract the 41-feature statistical representation described in Section IV. Traffic from the institution's own network is also captured during a 72-hour baseline collection period with all known intrusion sources excluded, providing a domain-specific normal-traffic baseline for the anomaly detection engine. Feature scaling using min-max normalisation is applied to ensure that features with large absolute ranges do not dominate distance-based model components.

The model training phase fine-tunes the Random Forest classifier using 5-fold stratified cross-validation on the combined training dataset, with hyperparameter optimisation via grid search over the number of trees (100–500),

maximum tree depth (10–30), and minimum samples per leaf (1–10). The Isolation Forest contamination parameter is set to 0.05, reflecting the expected proportion of anomalous traffic in normal operating conditions. Final model selection is based on the F1-score on the validation fold, which balances precision and recall across all attack categories.

System integration connects the trained models to the live packet processing pipeline through a Python extension module that accepts feature vectors from the C-implemented preprocessing layer via a shared memory interface, eliminating serialisation overhead at the detection boundary. The signature engine rule database is initialised with the Emerging Threats Open ruleset (approximately 27,000 rules) and supplemented with 150 locally authored rules targeting attack patterns specific to the institutional network environment.

The continuous refinement methodology schedules weekly retraining of the anomaly detection model on rolling 30-day windows of confirmed normal traffic, allowing the baseline to adapt to legitimate changes in network behaviour such as new application deployments or organisational restructuring. Newly confirmed intrusion samples are added to the labelled training set for the next scheduled retraining cycle, progressively improving detection coverage for attack patterns observed in the operational environment.

*Figure 2: NIDS Methodology Flow — Data Collection → Feature Extraction → Model Training → Signature Compilation → Hybrid Detection → Alert Generation → Analyst Review → Feedback Loop*

## VI. MODULES DESCRIPTION

The system is organized into seven functional modules, each encapsulating a distinct operational concern and exposing well-defined interfaces to adjacent modules.

### 6.1 Packet Capture Module

The packet capture module initialises one or more network interfaces in promiscuous mode using libpcap and begins continuous packet capture into a configurable ring buffer. The module supports Berkeley Packet Filter (BPF) expressions for pre-capture traffic filtering, allowing high-volume broadcast traffic to be excluded from the processing pipeline without consuming downstream resources. Captured packets are tagged with interface identifier, capture timestamp, and a sequence number for downstream ordering verification.

### 6.2 Protocol Decoder Module

The protocol decoder reconstructs application-layer sessions from the raw packet stream. TCP streams are reassembled in sequence-number order with configurable timeout for incomplete connections. The decoder identifies the application protocol from port numbers and protocol signatures (magic bytes) and invokes the corresponding decoder: HTTP, DNS, SMTP, FTP, SSH, or generic binary stream. Decoded session records contain the full flow metadata required for feature extraction.

### 6.3 Feature Extraction Module

For each completed session or time-windowed flow, the feature extraction module computes the 41-dimensional statistical feature vector. Features include: duration, protocol type, service, flag, source and destination byte counts, land indicator, wrong fragment count, urgent packet count, statistical counts of connections to the same host in the preceding two-second window (same host rate, same service rate, SYN error rate, REJ error rate), and similar two-second window statistics for connections to the same destination service. These features precisely replicate the NSL-KDD feature definition to ensure compatibility with models trained on that dataset.

### 6.4 Signature Matching Module

The signature matching module loads the compiled rule database at startup and applies the Aho-Corasick automaton to each packet payload. For rules with stateful constraints (e.g., threshold conditions requiring N occurrences within T seconds), the module maintains a per-rule state table implemented as a hash map keyed by source IP and destination IP-port pair. Rule updates trigger an atomic swap of the active and standby rule databases, ensuring zero-downtime updates with no detection gap during the swap window.

### 6.5 Anomaly Scoring Module

The anomaly scoring module receives the 41-dimensional feature vector and forwards it to the Random Forest classifier and Isolation Forest detector via the shared-memory interface. Each model returns a scalar anomaly score. The fusion function computes the weighted sum:  $\text{combined\_score} = 0.65 \times \text{RF\_score} + 0.35 \times \text{IF\_score}$ . Flows with a combined score exceeding the operational threshold (default 0.6) are forwarded to the alert generation module.

with an ANOMALY classification. Flows below the threshold are recorded in the flow statistics database for baseline monitoring.

## 6.6 Alert Management Module

The alert management module receives detection events from both the signature and anomaly engines, applies deduplication within a 60-second suppression window per source-destination pair per signature, and enriches each unique alert with contextual metadata. Alerts are assigned severity levels (Low, Medium, High, Critical) based on the signature priority or anomaly score magnitude. The module writes enriched alerts to an AES-256 encrypted SQLite alert database and forwards them to the dashboard via WebSocket. Analysts can acknowledge, escalate, suppress, or close alerts through the dashboard interface, with all actions logged immutably to the audit trail.

## 6.7 Reporting and Logging Module

The reporting module generates scheduled and on-demand reports in PDF and CSV formats covering alert volume trends, top attack sources, detection engine performance metrics, and false-positive rates estimated from analyst feedback. All log entries are written in structured JSON format to facilitate ingestion by external SIEM platforms such as Elasticsearch-Kibana or Splunk. Log files are rotated daily and encrypted before archival, with configurable retention periods to satisfy organisational and regulatory data retention requirements.

## VII. IMPLEMENTATION

The NIDS is implemented as a multi-process application on Ubuntu 22.04 LTS. The packet capture and signature detection components are implemented in C using libpcap 1.10.1 and the hyperscan 5.4 high-performance regular expression library, achieving packet processing rates exceeding 10 Gbps on commodity multi-core hardware. The anomaly detection engine is implemented in Python 3.10 using scikit-learn 1.2.0 and communicates with the C components via POSIX shared memory and UNIX domain sockets.

The management dashboard backend is implemented in Flask 2.2 with Flask-SocketIO for real-time WebSocket alert delivery. The frontend is implemented in React 18 with Recharts for traffic visualization and TailwindCSS for responsive layout. The alert database uses SQLite 3.39 with SQLCipher 4.5 for AES-256 encryption at rest. The complete system is containerised using Docker 24 with a docker-compose configuration that orchestrates the capture agent, detection engine, Flask API, and React frontend as separate services connected by an internal Docker network.

The training pipeline is implemented in a Jupyter notebook that can be executed independently of the live detection system. Model artefacts are serialised using Python pickle and loaded by the anomaly detection engine at startup. The Emerging Threats Open ruleset is downloaded and compiled into the Hyperscan database during container build, with a scheduled nightly rule update cron job that triggers a background reload without restarting the capture agent.

## VIII. TESTING AND RESULTS

The proposed NIDS was evaluated across four testing phases: controlled attack simulation using Metasploit and hping3, benchmark evaluation on the NSL-KDD test set, live campus network pilot deployment, and performance stress testing.

### 8.1 Controlled Attack Simulation

A testbed comprising five virtual machines (two attackers, one victim, one NIDS sensor, one management workstation) was configured on an isolated VLAN. Attack scenarios included port scanning (Nmap SYN scan, version detection scan), exploitation attempts (EternalBlue SMBv1 exploit, SSH brute force, HTTP directory traversal), denial-of-service (SYN flood, UDP flood, ICMP flood at 100,000 packets per second), and data exfiltration (DNS tunnelling, HTTPS covert channel). The NIDS correctly detected all 14 attack scenarios executed during testing, with alert generation latency below 800 milliseconds in all cases.

### 8.2 Benchmark Evaluation

On the NSL-KDD test set (22,544 instances, 4 attack categories: DoS, R2L, U2R, Probe), the hybrid detection engine achieved an overall accuracy of 97.3%, precision of 96.8%, recall of 97.1%, and F1-score of 96.9%. The signature engine alone achieved 94.1% on known attack signatures in the test set. The anomaly engine provided complementary coverage on the U2R (user-to-root) category, where signature matching achieved only 71.4% recall due to the low-volume, highly varied nature of privilege escalation attack patterns.

Table 1 compares the proposed hybrid NIDS against signature-only and anomaly-only baseline configurations on the NSL-KDD test set.

**Table 1: Detection Performance Comparison on NSL-KDD Test Set**

System Configuration	Accuracy	Precision	Recall	F1-Score
Signature-Only Baseline	91.4%	93.2%	89.7%	91.4%
Anomaly-Only (Random Forest)	94.8%	93.6%	95.9%	94.7%
Proposed Hybrid NIDS	97.3%	96.8%	97.1%	96.9%

### 8.3 Live Campus Network Pilot

The system was deployed on a monitoring span port of the campus core switch, processing approximately 2.3 Gbps of mixed traffic during peak hours. Over a 14-day pilot period, the system generated 1,847 unique alerts. Manual analyst review confirmed 1,721 true positives (93.2%), 89 false positives (4.8%), and 37 analyst-escalated ambiguous cases (2.0%). The false-positive rate was primarily attributable to the anomaly engine flagging unusual but legitimate scheduled data backup traffic between 02:00 and 04:00 daily; adding the backup server IP pair to the suppression list reduced the false-positive rate to 1.6%.

### 8.4 Performance Stress Testing

Packet processing throughput was measured under synthetic traffic loads generated by tcpreplay at rates from 1 Gbps to 40 Gbps on a server equipped with dual Intel Xeon Gold 6230 processors and a Mellanox ConnectX-5 100GbE NIC. The system sustained full signature and anomaly inspection at 12.4 Gbps without packet loss. At 15 Gbps and above, the anomaly engine became the throughput bottleneck due to per-flow feature computation latency; deploying an additional anomaly engine process on the second CPU socket extended the throughput ceiling to 24.8 Gbps.

**Table 2: System Performance Metrics**

Metric	Result
Overall Detection Accuracy (NSL-KDD)	97.3%
Signature Engine F1-Score	96.1%
Anomaly Engine F1-Score (Random Forest)	94.7%
Hybrid Engine F1-Score	96.9%
Alert Generation Latency (max)	< 800 ms
False-Positive Rate (post-tuning)	1.6%
Throughput (dual CPU socket)	24.8 Gbps
Pilot Deployment True Positive Rate	93.2%

## IX. CONCLUSION

The Secure Network Intrusion Detection System presented in this paper demonstrates that a hybrid signature-based and anomaly-based detection architecture can overcome the fundamental limitations of each approach individually, achieving 97.3% detection accuracy on the NSL-KDD benchmark while maintaining alert generation latency below 800 milliseconds and sustaining full-inspection throughput of up to 24.8 Gbps on commodity multi-socket server hardware.

The system's modular five-layer pipeline architecture — traffic capture, preprocessing and protocol analysis, parallel signature and anomaly detection engines, alert management, and management dashboard — provides clear separation of concerns that facilitates independent scaling, maintenance, and replacement of individual components. The encrypted alert database and immutable audit logging satisfy data integrity and confidentiality requirements essential for compliance with security governance frameworks.

The 14-day campus network pilot confirmed that the system is practically deployable in a real institutional environment, achieving a 93.2% true-positive rate that improved to a 98.4% true-positive rate after straightforward suppression tuning. The continuous retraining methodology ensures that the anomaly detection baseline adapts to legitimate network evolution without requiring manual baseline reconfiguration.

Future work will explore deep learning approaches — specifically transformer-based traffic sequence models and graph neural network representations of network flow graphs — to further improve detection of low-and-slow attack patterns that evade both signature matching and statistical anomaly detection. Integration with threat intelligence feeds for automated indicator-of-compromise enrichment and expansion of automated response capabilities to include DNS sinkholing and BGP blackhole routing are also planned as extensions to the current prototype.

## . REFERENCES

- [1] J. P. Anderson, “Computer Security Threat Monitoring and Surveillance,” Technical Report, James P. Anderson Co., Fort Washington, PA, 1980.
- [2] D. E. Denning, “An Intrusion-Detection Model,” *IEEE Trans. Softw. Eng.*, vol. SE-13, no. 2, pp. 222–232, Feb. 1987.
- [3] M. Tavallae, E. Bagheri, W. Lu, and A. A. Ghorbani, “A Detailed Analysis of the KDD CUP 99 Data Set,” in *Proc. IEEE Symp. Comput. Intell. Security Defense Appl.*, Ottawa, Canada, 2009, pp. 1–6.
- [4] I. Sharafaldin, A. H. Lashkari, and A. A. Ghorbani, “Toward Generating a New Intrusion Detection Dataset and Intrusion Traffic Characterization,” in *Proc. ICISSP*, Funchal, Portugal, 2018, pp. 108–116.
- [5] M. Roesch, “Snort: Lightweight Intrusion Detection for Networks,” in *Proc. USENIX LISA*, Seattle, WA, 1999, pp. 229–238.
- [6] A. L. Buczak and E. Guven, “A Survey of Data Mining and Machine Learning Methods for Cyber Security Intrusion Detection,” *IEEE Commun. Surv. Tutor.*, vol. 18, no. 2, pp. 1153–1176, Second Quarter 2016.
- [7] L. Breiman, “Random Forests,” *Mach. Learn.*, vol. 45, no. 1, pp. 5–32, Oct. 2001.
- [8] F. T. Liu, K. M. Ting, and Z.-H. Zhou, “Isolation Forest,” in *Proc. IEEE ICDM*, Pisa, Italy, 2008, pp. 413–422.
- [9] V. Chandola, A. Banerjee, and V. Kumar, “Anomaly Detection: A Survey,” *ACM Comput. Surv.*, vol. 41, no. 3, pp. 1–58, Jul. 2009.
- [10] Open Information Security Foundation, “Suricata IDS/IPS Documentation,” 2023. [Online]. Available: <https://suricata.io/documentation/>

