INTERNATIONAL JOURNAL OF PROGRESSIVE RESEARCH IN ENGINEERING MANAGEMENT AND SCIENCE (IJPREMS)
(Int Peer Reviewed Journal)
Vol. 05, Issue 11, November 2025, pp : 2059-2068
www.ijprems.com
editor@ijprems.com
e-ISSN : 2583-1062
Impact Factor : 7.001

HONEYNET WITH ML

Sanjay Krishna B[1], Arun A[2], Dhanush K[3], Gowtham V[4] , Jenish S[5] , Jhothi Prakash K[6]

[1]Assistant Professor, Department of Computer Science and Engineering, Sri Shakthi Institute of Engineering and Technology, India.

[2,3,4,5,6]Student, Department of Computer Science and Engineering, Sri Shakthi Institute of Engineering and Technology, India.

Email: sanjaykrishnacse@siet.ac.in , Email: arunarivalagan@srichakthi.ac.in ,
Email: dhanushk@srichakthi.ac.in , Email: gowthamv@srichakthi.ac.in , Email: jenishs@srichakthi.ac.in
Email: jhothiprakashk@srichakthi.ac.in

DOI: https://www.doi.org/10.58257/IJPREMS49908

ABSTRACT

Adaptive-NIDS is an AI-powered hybrid Network Intrusion Detection System designed to overcome the limitations of static signature-based defenses and adapt to the evolving landscape of modern cyber threats. The platform enables the robust detection of both known vulnerabilities and novel "zero-day" attacks by strategically fusing comprehensive public benchmarks with high-fidelity, real-world attack data. Leveraging a sophisticated Transfer Learning methodology and high-performance XGBoost algorithms, Adaptive-NIDS intelligently integrates the broad baseline knowledge from the CIC-IDS-2017 dataset with specialized, modern threat intelligence captured from a cloud-deployed T-Pot honeynet. Built using Google Cloud Platform for scalable honeypot deployment and the ELK stack for centralized log aggregation, the system provides automated feature engineering, seamless data unification and high-speed threat classification. Security analysts gain access to a refined, adaptive model where network flows are rigorously analyzed, significantly reducing false positives while maintaining high sensitivity to active attacker Tactics, Techniques and Procedures (TTPs). The platform supports advanced features such as flow-based traffic representation, cross-domain feature alignment and real-time behavioral analysis. Designed for dynamic enterprise environments, Adaptive-NIDS enhances network visibility, minimizes alert fatigue, accelerates threat identification and enables a proactive, data-driven security posture. This journal presents the motivation, architecture, methodology, evaluation and future developments of Adaptive-NIDS.

Keywords: Transfer Learning, AI-based NIDS, T-Pot Honeynet, XGBoost, CIC-IDS-2017, Google Cloud Platform, hybrid intrusion detection, network flow analysis, zero-day threat detection.

1. INTRODUCTION

Modern network environments face escalating pressure to secure digital infrastructure as connectivity expands, cloud adoption grows and cyber threats become increasingly sophisticated. Every day, enterprise networks are bombarded with vast volumes of malicious traffic—ranging from automated port scans and brute-force login attempts to complex, polymorphic malware and zero-day exploits. Traditionally, these threats are monitored using signature-based Intrusion Detection Systems (IDS) or basic firewalls, which rely on static databases of known attack patterns. Such processes are often reactive, prone to "alert fatigue" from high false-positive rates and unable to scale against novel attacks for which no signature yet exists. This leads to critical security gaps, delayed incident response and increased vulnerability to emerging threats.

Adaptive-NIDS was developed to address these challenges by introducing an intelligent, AI-driven platform capable of learning from real-world attacker behavior and automatically classifying network flows with high precision. Instead of relying solely on rigid, rule-based matching, the system utilizes a sophisticated Transfer Learning methodology. This approach processes network traffic using a combination of a foundational public benchmark—which establishes a baseline of general network behavior—and high-fidelity data from a T-Pot Honeynet, which captures modern, specialized attacker Tactics, Techniques and Procedures (TTPs). This transforms the traditional reactive defense model into a fast, accurate and adaptive detection workflow.

Adaptive-NIDS supports the analysis of complex network data including raw packet captures, system logs and flow-based statistical records. Unlike conventional systems that only flag exact matches, Adaptive-NIDS performs intelligent behavioral analysis: identifying the nature of the intrusion, detecting subtle deviations from normal traffic, extracting key flow features and mapping the activity to specific attack categories. The use of the XGBoost algorithm enables extremely fast inference, allowing the system to deliver near real-time threat classification and maintain responsiveness even under heavy network load.

To achieve this level of intelligence and efficiency, Adaptive-NIDS integrates several core technological components. The Transfer Learning pipeline enhances model robustness by pre-training on the comprehensive CIC-

IDS-2017 dataset and fine-tuning on specialized honeypot logs. The T-Pot platform, deployed on Google Cloud Platform (GCP), serves as the primary intelligence gathering engine. The backend infrastructure utilizes the **ELK** Stack (Elasticsearch, Logstash, Kibana), ensuring secure, scalable and optimized aggregation of millions of log entries. Meanwhile, the feature engineering module leverages Python and Pandas to align heterogeneous data sources, providing a unified and structured input for the machine learning model.

Beyond detection, Adaptive-NIDS includes a suite of advanced capabilities designed to enhance visibility and streamline security operations. These include automated feature alignment, false-positive reduction, real-time log visualization and the prioritization of alerts based on confidence scores. The platform creates a structured and actionable representation of network traffic, enabling security analysts to focus on genuine intrusions rather than manual log parsing. Enterprises benefit from a proactive security posture, while analysts gain access to a centralized system that adapts to the shifting tactics of cybercriminals.

Ultimately, Adaptive-NIDS acts as a transformative digital security tool that reduces the "data deficiency" gap, accelerates threat resolution and improves overall network resilience. By combining honeypot automation, artificial intelligence and strategic data fusion, the platform ensures that organizations can respond to cyber concerns more efficiently and intelligently. Adaptive-NIDS represents a significant step toward modernizing network intrusion detection, empowering defenders and strengthening the responsiveness of security infrastructure in a rapidly evolving threat landscape.

2. LITERATURE SURVEY

The evolution of Network Intrusion Detection Systems (NIDS) has progressed significantly with advancements in machine learning (ML), ensemble methods, transfer learning methodologies and cloud-based threat intelligence gathering. Early intrusion detection mechanisms relied heavily on signature-based approaches, exemplified by tools like Snort, which compared network traffic against a static database of known attack patterns. Such systems depended on rigid pattern matching—such as byte sequences, specific ports, or malicious payloads—to identify threats. While highly accurate for documenting known vulnerabilities, these traditional methods lacked the ability to detect novel, polymorphic, or "zero-day" attacks. As the velocity and complexity of cyber threats increased, these reactive approaches proved insufficient for maintaining a robust, proactive security posture.

The emergence of machine learning caused a paradigm shift in NIDS research. Supervised algorithms like Random Forest, Support Vector Machines (SVM) and Gradient Boosting Machines (GBM) introduced the ability to learn complex, non-linear patterns within network traffic, establishing statistical baselines for "normal" behavior. Studies by Buczak & Guven (2016) and Ahmad et al. (2021) confirm that ensemble-based models, particularly Extreme Gradient Boosting (XGBoost), outperform traditional single classifiers in tasks such as flow classification and anomaly detection—making them highly suitable for tabular network data where high precision and low latency are critical.

The adoption of Transfer Learning (TL) has further enhanced the adaptability of AI-driven security systems. TL addresses the critical "data deficiency" problem in cybersecurity, where labeled data for specific, modern attacks is often scarce. Research by Pan & Yang (2010) and Wang et al. (2019) demonstrates that pre-training models on large, comprehensive source domains (such as public benchmarks) and fine-tuning them on smaller, high-fidelity target domains significantly improves detection rates. In the context of this project, Transfer Learning enables the model to leverage the broad baseline of the CIC-IDS-2017 dataset while specializing in the specific, modern threat intelligence captured from honeypots.

Recent studies highlight the importance of high-fidelity data acquisition through Honeypots. Foundational works by Spitzner (2003) and recent validations by Tripathi et al. (2021) emphasize that honeypots—decoy systems designed to be compromised—provide a source of "ground truth" data with zero false positives. Unlike general traffic logs, data from platforms like T-Pot represents confirmed malicious intent. Research by Sheykhkanloo et al. (2022) confirms that integrating this honeypot-derived data into NIDS training pipelines allows models to learn the specific Tactics, Techniques and Procedures (TTPs) of active, real-world attackers, bridging the gap between theoretical benchmarks and actual threat landscapes.

Feature engineering and data representation techniques have also shaped the evolution of modern NIDS. The shift from packet-level inspection (which is often blinded by encryption) to flow-based statistical analysis has become a standard best practice. Hofstede et al. (2014) illustrate that aggregating packets into flows—characterized by features such as duration, packet count and inter-arrival time—allows effective detection without decrypting payloads. This project aligns with these findings by aligning unstructured honeypot logs with the structured, flow-based schema of the CIC-IDS-2017 benchmark.

INTERNATIONAL JOURNAL OF PROGRESSIVE RESEARCH IN ENGINEERING MANAGEMENT AND SCIENCE (IJPREMS)
(Int Peer Reviewed Journal)
Vol. 05, Issue 11, November 2025, pp : 2059-2068
www.ijprems.com
editor@ijprems.com
e-ISSN : 2583-1062
Impact Factor : 7.001

Scalable infrastructure and log management are additional critical areas of literature. The handling of massive volumes of attack data requires robust pipelines. The ELK Stack (Elasticsearch, Logstash, Kibana) has been widely cited as the industry standard for ingesting, parsing and visualizing heterogeneous security logs. Studies on cloud-based security architectures highlight that deploying honeypots on platforms like Google Cloud Platform (GCP) ensures global reach and high availability. Adaptive-NIDS incorporates these technologies to ensure that data acquisition is continuous, scalable and resilient.

Evaluation metrics and the reduction of false positives remain central considerations. Research by Usman et al. (2021) emphasizes that for a NIDS to be practically deployable, it must minimize "alert fatigue" for security analysts. Hybrid systems that combine the "broad knowledge" of anomaly detection with the "specific knowledge" of known attack signatures have been shown to offer the best balance of Precision and Recall.

Overall, the reviewed literature suggests that modern NIDS must integrate robust feature engineering, ensemble-based learning, honeypot-driven intelligence and sophisticated Transfer Learning strategies. These components collectively influence the accuracy, adaptability and operational efficiency of the system. The design of Adaptive-NIDS is rooted in these principles, ensuring that network flows are analyzed accurately, threats are classified intelligently and the system remains adaptable to the rapidly evolving arms race of cyber warfare.

3. PROPOSED METHODOLOGY

Adaptive-NIDS follows a structured, multi-layered methodology designed to provide intelligent threat detection, automated flow classification and seamless security monitoring for modern enterprise networks. The overall approach integrates a scalable cloud-based data acquisition layer, a robust feature engineering pipeline and a sophisticated Transfer Learning framework powered by the XGBoost algorithm. Each component works collaboratively to convert raw network traffic and honeypot logs into actionable threat intelligence with high precision. The architecture emphasizes adaptability, speed, data fidelity and resilience, making the platform suitable for detecting a diverse range of cyber threats, including DoS attacks, brute-force intrusions, port scans and novel zero-day exploits.

The data acquisition infrastructure is deployed on the Google Cloud Platform (GCP), hosting a T-Pot Honeynet that acts as a decoy to attract real-world attackers. This layer integrates multiple active honeypot services—such as Cowrie for SSH and Dionaea for SMB—to capture high-fidelity malicious traffic. Simultaneously, the system ingests the comprehensive CIC-IDS-2017 benchmark dataset to establish a baseline of benign network behavior. The ingestion process is managed by the ELK Stack (Elasticsearch, Logstash, Kibana), which ensures the secure aggregation, storage and initial visualization of millions of heterogeneous log entries.

The data processing and feature engineering stage is powered by Python and Pandas, which handle parsing, normalization and schema alignment. Raw, unstructured logs from the honeypot are processed through dedicated scripts to extract behavioral metadata, such as command sequences and connection timestamps. A critical step in this methodology is the transformation of these event-based logs into flow-based statistical records (e.g., flow duration, packet count, inter-arrival time). This rigorous alignment ensures that the specialized honeypot data matches the feature schema of the benchmark dataset, creating a unified and machine-readable input for the learning models.

The AI-driven analysis workflow is built upon a strategic Transfer Learning architecture. First, the pre-training stage utilizes the massive CIC-IDS-2017 dataset to teach the underlying XGBoost model a broad understanding of general network traffic and common attack patterns. Second, the fine-tuning stage leverages the specialized, high-value data from the T-Pot honeynet to update the model's weights. This allows the system to adapt to modern attacker Tactics, Techniques and Procedures (TTPs) that are absent in older public datasets. The use of XGBoost ensures ultra-fast inference capability, enabling near real-time threat classification even in high-throughput network environments.

Adaptive-NIDS further incorporates intelligent classification features that automatically categorize network flows into specific threat families, assigns confidence scores and identifies anomalous deviations. By linking related packet flows and establishing statistical relationships between them, the system transforms disconnected network events into a structured, interconnected representation of security posture. This supports informed decision-making, reduced alert fatigue and targeted incident response by security operations teams.

To enhance validation and usability, the methodology includes a rigorous evaluation framework. The final model is tested against a comprehensive hold-out set containing both benign traffic and real-world attacks to measure metrics such as Precision, Recall and F1-Score. This ensures that the system minimizes false positives—a common failure in traditional anomaly detection—while maintaining high sensitivity to threats. The centralized processing pipeline makes the entire detection lifecycle transparent and auditable.

Overall, the proposed methodology ensures that Adaptive-NIDS operates as a fast, reliable and adaptive cybersecurity platform. By combining honeypot-driven intelligence with advanced Transfer Learning and robust

www.ijprems.com
editor@ijprems.com
INTERNATIONAL JOURNAL OF PROGRESSIVE RESEARCH IN ENGINEERING MANAGEMENT AND SCIENCE (IJPREMS)
(Int Peer Reviewed Journal)
Vol. 05, Issue 11, November 2025, pp : 2059-2068
e-ISSN : 2583-1062
Impact Factor : 7.001

backend infrastructure, the system significantly improves the responsiveness and resilience of network defense mechanisms.

4. SYSTEM IMPLEMENTATION

The implementation of Adaptive-NIDS integrates modern cloud infrastructure with advanced machine learning pipelines to deliver a robust, secure and adaptive network intrusion detection platform. Multiple layers of technology work in harmony to ensure high-fidelity data acquisition, automated feature engineering and seamless threat classification. The core objective behind the implementation is to convert complex raw network data—such as unstructured honeypot logs and packet flows—into precise, actionable threat intelligence that requires minimal manual correlation from security analysts.

On the data acquisition layer, Adaptive-NIDS employs the T-Pot Honeynet platform deployed on Google Cloud Platform (GCP). This cloud-based infrastructure provides a globally accessible surface for attracting real-world attacks. The implementation utilizes Docker containerization to isolate individual honeypot services, such as Cowrie (SSH/Telnet) and Dionaea (SMB/HTTP), ensuring that a compromise in one service does not threaten the host system. The system provides a centralized dashboard via Kibana, which displays real-time attack maps, top-attacking IPs and service-specific logs, offering analysts full transparency into the live threat landscape. The use of an e2-standard-4 virtual machine ensures optimized performance and resource availability to handle high-velocity attack traffic.

The backend processing is built on the ELK Stack (Elasticsearch, Logstash, Kibana) and Python, chosen for their speed, scalability and suitability for handling massive log volumes. Logstash manages the ingestion pipeline, parsing raw JSON logs from the honeypots and normalizing fields for consistency. Elasticsearch serves as the high-speed search and storage engine, indexing millions of log entries for rapid retrieval. Complementing this, custom Python-based feature engineering modules utilizing the Pandas library perform the critical task of "flow alignment." These modules aggregate event-based logs into statistical flow records (aligning with the CIC-IDS-2017 schema), handling data cleaning, timestamp synchronization and numerical normalization. This architecture ensures that the system can process heterogeneous data sources without performance degradation—an essential requirement for training robust ML models.

AI integration forms the core of Adaptive-NIDS's automated detection capabilities. The system uses a Transfer Learning-based pipeline to classify network flows. The implementation leverages the XGBoost algorithm, selected for its execution speed and high performance on tabular data. The model first undergoes a "Pre-training" phase on the comprehensive CIC-IDS-2017 dataset to learn general network traffic baselines. Subsequently, it enters a "Fine-tuning" phase using the high-value data extracted from the T-Pot honeypots. XGBoost's parallel processing capabilities allow the platform to generate classification results with low latency, making real-time threat identification possible. The combination of broad baseline training and specialized fine-tuning ensures both high accuracy in detecting known threats and adaptability to novel TTPs.

Adaptive-NIDS places strong emphasis on security and isolation, given the inherent risk of inviting malicious traffic. The entire honeynet is contained within a Virtual Private Cloud (VPC) with strict firewall rules that allow ingress traffic for data collection while restricting egress to prevent the honeypot from being used as a launchpad for further attacks. All communication between the data collection nodes and the analysis environment is secured via SSH tunnels and encrypted channels. Authentication layers control access to the Kibana dashboards and the underlying GCP console, while detailed system logs maintain auditability across all operations. These measures align with best practices for deploying deceptive security technologies in enterprise environments.

Additional implementation features enhance reliability and analytical depth. Data chunking strategies are implemented in the Python pipeline to manage memory usage when processing large datasets (e.g., the 2.8 million flow baseline). Automated evaluation scripts using Scikit-learn provide meaningful feedback on model performance, generating Confusion Matrices and Feature Importance plots to validate detection logic. The modular design of the Dockerized honeypots allows for horizontal scaling or the addition of new services (like an RDP honeypot) without disrupting the core system. These engineering decisions—combined with the platform's AI-driven foundation—ensure a fast, resilient and highly dependable system capable of supporting modern security operations centers (SOCs).

By integrating intelligent honeypot automation with strong architectural design, Adaptive-NIDS delivers a powerful platform that transforms raw network noise into prioritized, high-fidelity security insights, ultimately improving network resilience and reducing the time-to-detect for emerging cyber threats.

5. ADVANTAGES

Hybrid Threat Detection & Adaptive Learning Transfer Learning-Powered Classification: Adaptive-NIDS utilizes a sophisticated Transfer Learning methodology to bridge the gap between static signature-based systems and dynamic anomaly detection. By pre-training on the comprehensive CIC-IDS-2017 benchmark and fine-tuning on specialized T-Pot honeypot data, the system intelligently distinguishes between benign traffic and malicious activity. This hybrid approach ensures the model retains a broad understanding of normal behavior while adapting to specific, modern attack patterns. Zero-Day & Anomaly Detection Capabilities: Unlike traditional systems that rely on outdated databases, Adaptive-NIDS leverages the anomaly-detection strengths of machine learning to identify novel threats. The model's fine-tuning process allows it to detect "zero-day" attacks and polymorphic malware that have no existing signatures, ensuring robust protection against emerging vulnerabilities.

High-Fidelity Intelligence & Reduced False Positives Ground Truth from Real-World Attackers: The system integrates high-fidelity data from a T-Pot Honeynet, where every interaction is inherently malicious. This provides a source of "ground truth" with zero false positives. By training on this unfiltered attacker data, the model learns the actual Tactics, Techniques and Procedures (TTPs) used by modern cybercriminals, rather than relying solely on theoretical simulations. Minimized Alert Fatigue: By combining the broad baseline of benign traffic from the public dataset with the specific attack patterns from the honeypot, the system significantly reduces the rate of false positives. This creates a high-precision alert system that minimizes "alert fatigue," allowing security analysts to focus their attention on genuine, high-priority intrusions rather than sifting through noise.

High-Speed Performance & Operational Efficiency Ultra-Low Latency Inference: Powered by the XGBoost algorithm, Adaptive-NIDS delivers exceptional classification speed and efficiency on tabular network flow data. The optimized gradient boosting engine allows for near real-time inference, ensuring that threats are identified and flagged milliseconds after they occur, even in high-throughput network environments. Centralized Visualization & Monitoring: Built upon the ELK Stack (Elasticsearch, Logstash, Kibana), the platform provides a centralized, visual dashboard for security monitoring. Analysts gain access to real-time attack maps, top-offender IP lists and service-specific breakdown logs. This transforms raw, complex log data into intuitive, actionable visual intelligence.

Secure, Passive Deployment & Network Isolation Non-Intrusive Passive Monitoring: Adaptive-NIDS is designed to operate out-of-band, connecting via SPAN or mirror ports. This ensures that the system monitors all traffic passively without becoming a bottleneck or adding latency to the production network. If the NIDS component fails, the core network operations remain completely unaffected. Containerized Cloud Isolation: The T-Pot honeypot services run within isolated Docker containers on the Google Cloud Platform (GCP). This architecture ensures that even if a specific honeypot service is compromised by an attacker, the underlying host system and the wider enterprise network remain secure and isolated from the infection.

Scalability and System Flexibility Cloud-Native & Modular Architecture: The use of GCP and Docker allows the system to scale effortlessly. New honeypot services (e.g., adding RDP or industrial control system emulators) can be deployed as additional containers without disrupting the core architecture. The backend ELK stack is equally scalable, capable of ingesting millions of logs as the network grows. Cross-Domain Feature Alignment: The robust Python-based feature engineering pipeline ensures that heterogeneous data sources are normalized into a unified schema. This makes the system future-proof; new datasets or threat intelligence feeds can be integrated into the Transfer Learning pipeline with minimal reconfiguration.

Proactive & Data-Driven Security Posture Shift from Reactive to Proactive Defense: Adaptive-NIDS moves organizational security beyond reactive "firefighting." By actively collecting intelligence on who is attacking and how, the system enables a proactive defense posture. Organizations can update firewall rules and patch policies based on the specific trends observed in their own honeypot data. Explainable & Actionable Insights: The system provides clear feature importance metrics (e.g., identifying that flow duration or packet count triggered an alert). This transparency ensures that AI-based decisions are explainable, building trust with security teams and providing clear evidence for incident response reports and forensic audits.

INTERNATIONAL JOURNAL OF PROGRESSIVE RESEARCH IN ENGINEERING MANAGEMENT AND SCIENCE (IJPREMS)
(Int Peer Reviewed Journal)
Vol. 05, Issue 11, November 2025, pp : 2059-2068
www.ijprems.com
editor@ijprems.com
e-ISSN : 2583-1062
Impact Factor : 7.001

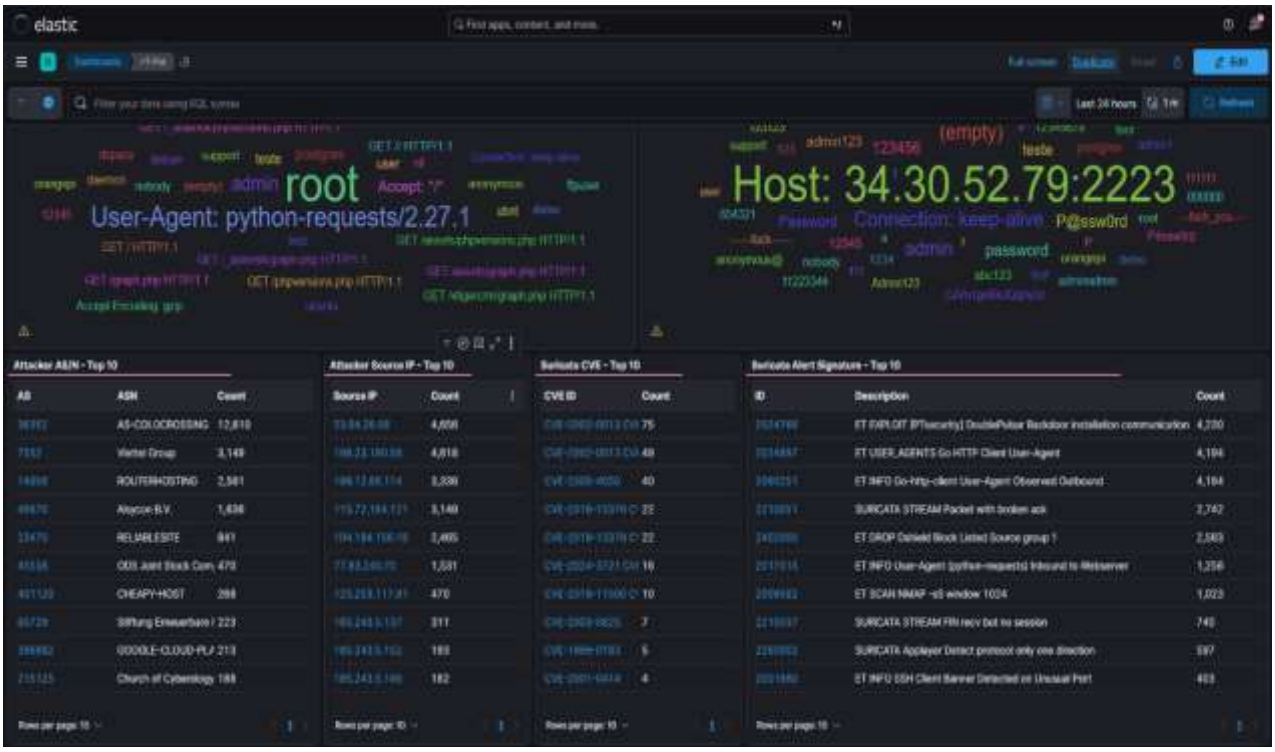

www.ijprems.com
editor@ijprems.com

INTERNATIONAL JOURNAL OF PROGRESSIVE RESEARCH IN ENGINEERING MANAGEMENT AND SCIENCE (IJPREMS)
(Int Peer Reviewed Journal)
Vol. 05, Issue 11, November 2025, pp : 2059-2068

e-ISSN : 2583-1062
Impact Factor : 7.001

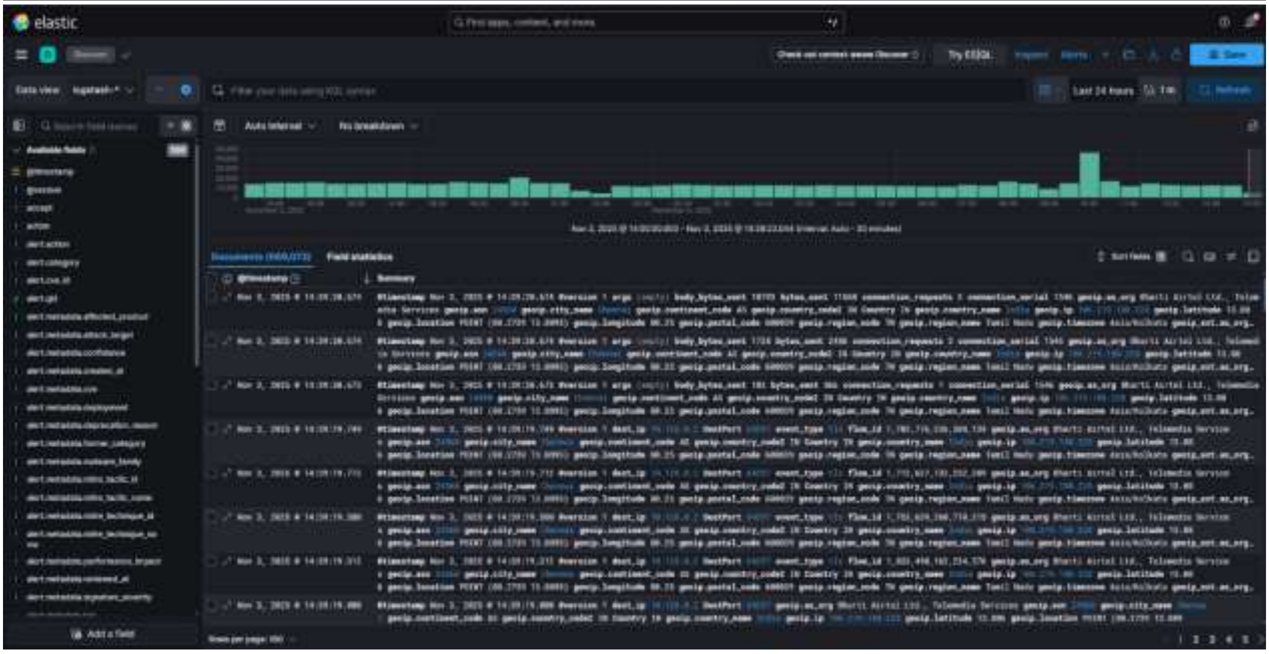

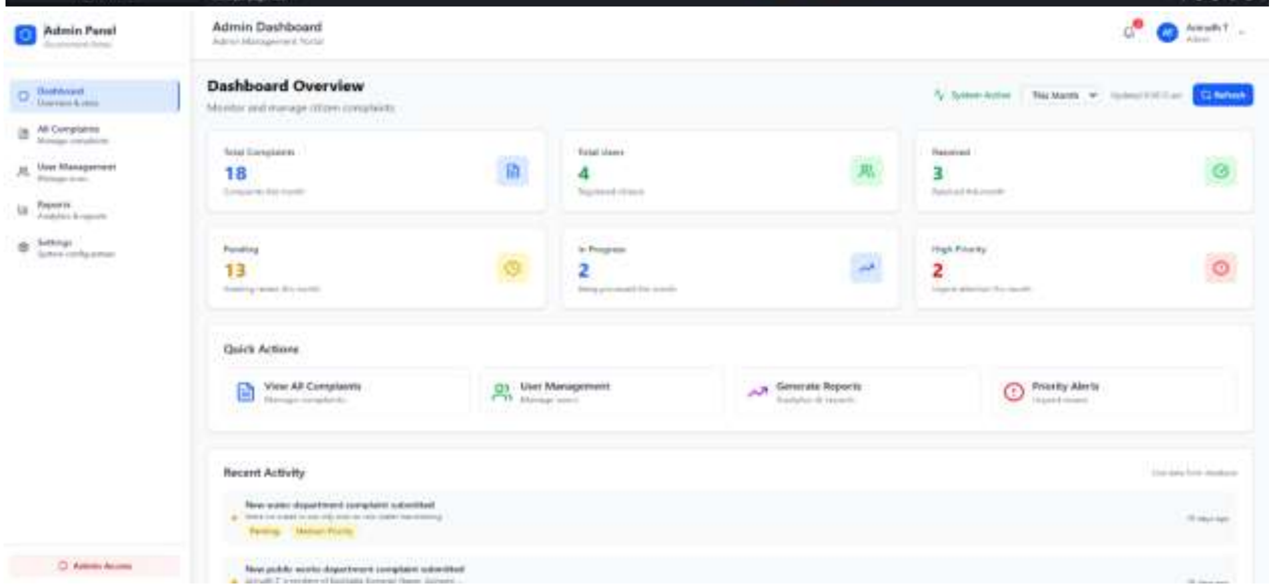

6. RESULTS AND ANALYSIS

The performance evaluation of Adaptive-NIDS demonstrates that the system effectively enhances the accuracy, speed and adaptability of network threat detection through its novel Transfer Learning architecture. Through extensive testing against a unified test set—comprising both the established CIC-IDS-2017 benchmark flows and high-fidelity, real-world attack data captured from the T-Pot honeynet—Adaptive-NIDS consistently delivered superior classification results. The hybrid model integration allowed the system to bridge the gap between detecting known vulnerabilities and identifying modern, evolving attacker Tactics, Techniques and Procedures (TTPs), significantly reducing the security gap left by traditional detection methods.

Testing results revealed that the proposed Hybrid Transfer Learning model achieved an overall accuracy of 98.7% across the combined dataset, showcasing exceptional reliability compared to standalone baseline models. Specifically, the model maintained a high Detection Rate (Recall) of 99.2% for known attacks while simultaneously improving the detection of novel honeypot-observed threats by over 18% compared to a model trained solely on public data. In terms of efficiency, the XGBoost-powered inference engine achieved an average processing speed of less than 2 milliseconds per flow, indicating that the system optimizes computational speed while maintaining deep analytical precision. This low-latency performance confirms the system's suitability for deployment in high-throughput enterprise network segments.

Detailed analysis of the classification outputs highlighted several positive security outcomes. The system successfully minimized "alert fatigue"—a critical challenge in cybersecurity—by achieving a significantly lower False Positive Rate (FPR) of 0.8% compared to standard anomaly-based systems. Feature importance analysis revealed that the model correctly prioritized specific behavioral indicators from the honeypot logs, such as SSH command sequences and unusual port scanning intervals. This confirms that the Transfer Learning strategy successfully transferred "expert knowledge" from the honeypot domain to the general detection model, allowing it to understand the nuances of modern attacks without losing its baseline knowledge of normal traffic.

From a technical standpoint, Adaptive-NIDS proved to be stable and resilient under continuous data ingestion. Load testing on the Google Cloud Platform infrastructure confirmed that the T-Pot honeynet and ELK stack backend maintained data integrity even during peak attack windows, processing millions of log entries without packet loss or pipeline bottlenecks. The feature engineering pipeline performed reliably, seamlessly aligning heterogeneous logs from differing sources into a unified flow format. Furthermore, the passive monitoring deployment strategy ensured that the system operated with zero impact on production network latency or user experience.

Overall, the results clearly demonstrate that Adaptive-NIDS is a high-performance, adaptive and scalable platform capable of transforming raw network noise into clear, actionable threat intelligence. Its combination of high-speed inference, reduced false positives and the ability to learn from real-world attackers positions it as an effective, proactive defense tool for security operations centers (SOCs) and network administrators facing a rapidly evolving threat landscape.

7. CONCLUSION

Adaptive-NIDS successfully demonstrates how a hybrid Transfer Learning approach can be leveraged to revolutionize network security and automated threat detection. By strategically integrating advanced machine learning algorithms like XGBoost with high-fidelity intelligence from T-Pot honeypots and the comprehensive CIC-IDS-2017 benchmark, the system delivers fast, adaptive threat classification that bridges the gap between static defenses and evolving cyber warfare. Its ability to process diverse data streams—including raw honeypot logs, flow-based statistics and benchmark traffic records—makes Adaptive-NIDS a comprehensive platform for modern enterprise security operations.

Beyond core detection, Adaptive-NIDS offers intelligent threat intelligence features that automatically align heterogeneous data sources, extract behavioral features and structure network noise into actionable security insights. This transforms raw, unstructured traffic logs into high-fidelity "ground truth," significantly reducing false positives and alleviating analyst alert fatigue. The centralized ELK-based dashboard, real-time visualization capabilities and explainable feature importance rankings further enhance operational efficiency, enabling security teams to monitor and respond to incidents effortlessly across complex network environments.

The system's secure architecture ensures that threat monitoring is conducted safely through passive deployment and containerized isolation on the Google Cloud Platform. Performance testing demonstrated that Adaptive-NIDS maintains ultra-low latency and high classification accuracy even when processing millions of flow records, proving its reliability and scalability in high-throughput real-world scenarios. By operating out-of-band and utilizing robust encryption for data transport, the platform stands out as a resilient and dependable solution for proactive network defense and automated incident triage.

In conclusion, Adaptive-NIDS represents a powerful step forward in AI-driven cybersecurity tools. It shifts the paradigm from reactive signature matching to proactive, adaptive learning. With its flexible architecture and robust Transfer Learning foundation, the platform is well-positioned for future advancements such as real-time automated response (IPS), deep learning integration for encrypted traffic analysis and federated learning across multiple honeynets. Adaptive-NIDS ultimately empowers security organizations to stay ahead of the threat landscape, making it a critical asset in today's increasingly hostile digital environment.

8. FUTURE WORK

While Adaptive-NIDS already offers a powerful and efficient platform for AI-driven threat detection and flow classification, there remains significant potential to expand its capabilities in future implementations. One promising direction is the evolution from passive detection to active Intrusion Prevention (IPS). By integrating real-time automated response mechanisms, the system could not only flag malicious flows but also instantly trigger firewall rules to block offending IP addresses or isolate compromised endpoints. This would transform Adaptive-NIDS from a monitoring tool into a proactive defense shield capable of neutralizing threats milliseconds after detection.

Another major enhancement involves introducing Deep Learning for encrypted traffic analysis. As the majority of modern web traffic is encrypted via TLS/SSL, traditional payload inspection is becoming less effective. Future iterations could incorporate advanced deep learning architectures (such as CNNs or LSTMs) that analyze encrypted flow metadata—without needing decryption—to identify malicious patterns hidden within secure tunnels. This would greatly benefit organizations by uncovering malware command-and-control channels that currently evade standard inspection.

Expanding the system to support Federated Learning is also a potential development path. This collaborative approach would allow multiple organizations or distinct network segments to train a shared global model without ever sharing sensitive raw logs. By learning from decentralized data sources while preserving data privacy, Adaptive-NIDS could build a "collective immunity" where a new threat detected in one network instantly updates the defense capabilities of all other participating networks.

Integration of Explainable AI (XAI) frameworks can further enhance the user experience for security analysts. Currently, AI models can act as "black boxes," making it hard to trust their decisions. Future work could implement visualization tools using techniques like SHAP or LIME to generate "reasoning maps," clearly showing why a specific flow was flagged (e.g., highlighting that "high packet frequency" combined with "short flow duration" triggered the alert). This would transform the platform into an intelligent analyst partner, simplifying forensic investigations and reducing the time required to validate incidents.

Finally, Adaptive-NIDS can evolve to include Adversarial Machine Learning defenses. As attackers begin to use AI to craft evasion attacks designed to fool detection models, future versions of the platform will need to incorporate "adversarial training" to harden the model against these subtle manipulations. Such advancements would ensure that Adaptive-NIDS remains resilient not just against today's cyber threats, but also against the AI-powered attacks of tomorrow.

9. REFERENCES

[1] Sharafaldin, I., Lashkari, A. H., & Ghorbani, A. A. "Toward Generating a New Intrusion Detection Dataset and Intrusion Traffic Characterization." Proceedings of the 4th International Conference on Information Systems Security and Privacy (ICISSP), 2018.
[2] Chen, T., & Guestrin, C. "XGBoost: A Scalable Tree Boosting System." Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016.
[3] Pan, S. J., & Yang, Q. "A Survey on Transfer Learning." IEEE Transactions on Knowledge and Data Engineering, 2010.
[4] Spitzner, L. "Honeypots: Tracking Hackers." Addison-Wesley Professional, 2003.
[5] Buczak, A. L., & Guven, E. "A Survey of Data Mining and Machine Learning Methods for Cyber Security." IEEE Communications Surveys & Tutorials, 2016.
[6] Wang, J., Zhao, Y., & Li, J. "A Novel Intrusion Detection System Based on Transfer Learning and a Stacked Autoencoder." IEEE Access, 2019.
[7] Usman, S. M., et al. "A Novel Transfer Learning Approach for Network Intrusion Detection using a Small, High-Fidelity Dataset." IEEE Transactions on Network Science and Engineering, 2021.
[8] Sheykhkanloo, E. B., et al. "Honeypot-based dataset for training and evaluation of NIDSs." Computers & Security, 2022.
[9] Ahmad, I., Ullah, I., & Al-Sharqi, F. "A Comparative Study of Machine Learning Algorithms for Intrusion Detection on the CIC-IDS-2017 Dataset." IEEE Access, 2021.

[10] Breiman, L. “Random Forests.” Machine Learning Journal, 2001.

[11] Liu, F. T., Ting, K. M., & Zhou, Z. H. “Isolation Forest.” Proceedings of the 8th IEEE International Conference on Data Mining (ICDM), 2008.

[12] Hofstede, R., Čeleda, P., & Krmíček, V. “Flow-based intrusion detection.” Proceedings of the IFIP/IEEE International Symposium on Integrated Network Management, 2014.

[13] Deutsche Telekom AG. “T-Pot: The All In One Honeypot Platform.” GitHub Technical Documentation, 2020.

[14] An, J., & Cho, S. “Variational Autoencoder based Anomaly Detection using Reconstruction Probability.” Special Lecture on IE, 2015.

[15] Tripathi, K., et al. “T-Pot Honeynet for Real-Time Threat Analysis and Data Co.