AI-Driven Next-Generation Firewall for Dynamic Threat Detection and Zero Trust Implementation

Shivam Kumar, Hanshika Shanvi, Randhir Kumar, Santosh Kumar, Mr. Deepesh Kumar, Mr. Badal Bhushan

Department of Computer Science and Engineering, IIMT College of Engineering, Greater Noida, Uttar Pradesh, India

DOI: https://doi.org/10.51584/IJRIAS.2025.10120052

Received: 26 December 2025; Accepted: 31 December 2025; Published: 15 January 2026

ABSTRACT

The increasing adoption of cloud computing, remote work environments, Internet of Things (IoT) devices, and encrypted communication has significantly expanded the attack surface of modern enterprise networks. Traditional rule-based and signature-driven firewall systems are no longer sufficient to defend against advanced cyber threats such as zero-day attacks, lateral movement, and stealthy intrusion attempts. These conventional approaches lack adaptability, generate high false-positive rates, and fail to provide continuous trust evaluation required in dynamic network environments.

To address these limitations, this paper proposes an AI-driven Next-generation firewall (NGFW) architecture designed to support dynamic threat detection and Zero Trust implementation. The proposed framework integrates network traffic monitoring, behavioral flow analysis, AI-based threat detection, and dynamic policy enforcement into a unified security system. By analyzing traffic patterns at the flow level, the system continuously evaluates risk and enforces least-privilege access decisions without relying on static rules or predefined signatures.

The effectiveness of the proposed approach is evaluated using publicly available intrusion detection datasets and a hybrid enterprise network testbed. Experimental results demonstrate that the AI-driven NGFW achieves higher detection accuracy, reduced false-positive rates, and faster policy adaptation compared to traditional rule-based and signature-based security solutions. These findings validate that integrating AI-based traffic analysis with Zero Trust principles significantly enhances network security, adaptability, and operational efficiency in modern enterprise environments.

Keywords— Next-generation firewall (NGFW), AI-Driven Network Security, Dynamic Threat Detection, Zero Trust architecture (ZTA), Behavioral Traffic Analysis, Intrusion Detection, Adaptive Security Policy Enforcement

INTRODUCTION

Modern enterprise networks are undergoing rapid transformation due to the widespread adoption of cloud services, remote access technologies, Internet of Things (IoT) devices, and software-defined infrastructures. While these advancements improve scalability and flexibility, they also introduce complex security challenges by expanding network boundaries and increasing exposure to sophisticated cyber threats. Attacks such as zero-day exploits, advanced persistent threats (APTs), insider misuse, and lateral movement across internal networks pose serious risks to traditional security mechanisms.

Conventional firewalls and intrusion detection systems (IDS/IPS) primarily rely on static rule sets and predefined attack signatures. Although effective against known threats, these systems struggle to detect evolving and previously unseen attacks. Additionally, encrypted traffic and context-dependent attack behaviors reduce the effectiveness of deep packet inspection, creating blind spots in network visibility. As a result, security teams often face excessive false positives, delayed threat detection, and limited situational awareness.

Implementing Zero Trust at scale requires intelligent and automated security mechanisms capable of continuously analyzing network behavior and dynamically enforcing access policies. Despite advances in artificial intelligence for cybersecurity, existing research often treats threat detection, policy enforcement, and Zero Trust implementation as separate components. There is a lack of integrated frameworks that combine AI-based behavioral threat detection with dynamic Zero Trust policy enforcement within a single NGFW system.

This paper addresses this gap by proposing an AI-driven Next-generation firewall architecture that enables dynamic threat detection, continuous trust evaluation, and adaptive policy enforcement aligned with Zero Trust principles. The proposed system analyzes network traffic at the flow level using AI-based behavioral modeling to identify malicious activities and enforce real-time security decisions.

This paper proposes an AI-driven Next-generation firewall based on behavioral traffic analysis and machine learning techniques for dynamic threat detection, malicious flow identification, and Zero Trust policy enforcement.

1) Design of an AI-driven NGFW architecture capable of detecting advanced and dynamic network threats.

2) Integration of behavioral traffic analysis for accurate detection and localization of malicious network flows.

3) Dynamic Zero Trust policy enforcement based on continuous risk assessment rather than static access rules.

The rest of this paper is organized as follows. Section II reviews related work on AI-based next-generation firewalls and Zero Trust security models. Section III describes the proposed AI-driven firewall architecture and threat detection methodology. Section IV presents the experimental setup and performance evaluation. Section V discusses the results and challenges. Finally, Section VI concludes the paper and outlines future research directions.

Related Work

Network security has traditionally relied on perimeter-based firewalls and signature-driven intrusion detection systems to protect enterprise environments. These mechanisms primarily depend on static rules and known attack signatures, which limits their ability to detect zero-day attacks, encrypted traffic anomalies, and evolving threat behaviors. As modern networks become more dynamic and distributed, the effectiveness of such traditional security approaches has significantly declined.

To overcome these limitations, Next-generation firewalls (NGFWs) were introduced, integrating features such as application-level inspection, intrusion prevention, and traffic monitoring. NGFWs provide improved visibility compared to traditional firewalls; however, many existing solutions still rely on predefined policies and heuristic-based detection techniques. As a result, they face challenges in adapting to rapidly changing attack patterns and large-scale network environments.

Parallel to advancements in AI-based threat detection, the Zero Trust architecture (ZTA) has emerged as a prominent security model for modern enterprises. Zero Trust eliminates implicit trust assumptions and enforces continuous verification, least-privilege access, and contextual decision-making for every network interaction. Existing studies highlight the importance of intelligent policy enforcement mechanisms to support Zero Trust; however, many implementations rely on static access rules and lack adaptive threat intelligence.

Overall, existing literature indicates a gap in unified frameworks that combine AI-driven dynamic threat detection, explainable security decisions, and Zero Trust policy enforcement within a Next-generation firewall system. This paper addresses this gap by presenting an AI-driven NGFW approach that integrates adaptive threat analysis with Zero Trust principles to enhance network security effectiveness in modern enterprise environments.

PROPOSED METHODOLOGY

This section describes the methodology adopted to design and evaluate an AI-driven Next-generation firewall (NGFW) framework that supports dynamic threat detection and Zero Trust implementation. The proposed approach integrates network traffic analysis, AI-assisted threat detection, explainable decision-making, and policy enforcement into a unified security framework.

System Overview

The proposed methodology follows a modular architecture consisting of traffic collection, AI-based analysis, decision interpretation, and Zero Trust policy enforcement. Network traffic is continuously monitored and analyzed to identify potential security threats. The framework emphasizes adaptability, transparency, and continuous verification, which are essential requirements for modern enterprise security systems.

Network Traffic Data Collection

Network traffic data is collected from enterprise network environments using standard monitoring mechanisms such as packet capture and flow logs. The collected data includes network flow attributes such as source and destination addresses, protocol types, session duration, and traffic volume. These features provide behavioral context necessary for identifying abnormal and malicious activities.

To support evaluation, publicly available network intrusion datasets are utilized. These datasets contain labeled traffic instances representing normal behavior and various attack categories, enabling systematic performance analysis.

AI-Based Threat Detection Module

The threat detection module employs artificial intelligence techniques to analyze network traffic behavior. The AI model is trained to distinguish between benign and malicious traffic based on learned patterns. Unlike rule-based firewalls, the AI-driven approach adapts to evolving traffic characteristics and supports detection of previously unseen threats.

The detection process operates at the network flow level, allowing the system to localize suspicious sessions and identify potential intrusion points. This behavior-based analysis enhances detection accuracy while reducing reliance on static signatures.

Machine Learning Model Specification

The AI-driven threat detection module employs a hybrid learning strategy combining supervised and unsupervised machine learning techniques to address both known and previously unseen threats. For supervised classification of known attack categories, a Random Forest classifier is utilized due to its robustness against noisy features, scalability, and interpretability. To detect unknown or zero-day threats, an Autoencoder-based unsupervised anomaly detection model is used to learn normal traffic behavior and identify deviations.

The models are trained on flow-level features extracted from network traffic, including packet count, byte volume, inter-arrival time, session duration, protocol distribution, and traffic directionality. Training is performed using an 80:20 train–test split with five-fold cross-validation to ensure generalization and robustness. Hyperparameters are optimized through grid search to improve detection accuracy while minimizing false positives.

Dynamic Threat Adaptation

To address dynamic and evolving threats, the proposed framework supports continuous learning and contextual analysis. Traffic patterns are evaluated over time to capture changes in behavior that may indicate emerging attacks. This adaptive mechanism enables the firewall to respond to new threat patterns without requiring manual rule updates.

Explainable Security Decision Layer

Explainability is incorporated into the decision-making process to improve transparency and operational trust. For each detected threat, the system generates interpretable explanations describing the factors that contributed to the detection decision. These explanations assist security analysts in understanding the nature of detected threats and facilitate timely response actions.

The explainable layer supports improved incident analysis and reduces alert fatigue by providing meaningful contextual information alongside security alerts.

Example of Explainable Threat Decision

To demonstrate explainability, consider a detected lateral movement attempt within the enterprise network. The explainable decision layer identified abnormal session duration (32%), access to sensitive internal assets (27%), and protocol misuse (21%) as the primary contributors to the threat classification. Based on these factors, the session risk score exceeded the Zero Trust threshold, resulting in immediate access restriction and policy escalation.

This interpretable output enables security analysts to understand why a session was flagged as malicious, improving trust, auditability, and incident response efficiency.

Zero Trust Policy Enforcement

The proposed framework integrates Zero Trust principles by enforcing continuous verification and least-privilege access control. Access decisions are made based on contextual risk assessment derived from AI-based threat analysis. Network entities are granted access only when they satisfy predefined security conditions, and access privileges are continuously re-evaluated during active sessions.

This dynamic policy enforcement mechanism minimizes implicit trust and strengthens network security by limiting unauthorized or suspicious activities.

In a simulated enterprise scenario, an authenticated employee attempted to access a restricted database from an unusual location and device. The AI-driven firewall detected abnormal access behavior and elevated the session risk score. As a result, the Zero Trust policy engine enforced adaptive authentication and restricted access until verification was completed, demonstrating real-time Zero Trust enforcement in practice.

System Architecture

Although the proposed architecture follows a layered AI pipeline, it is specifically tailored for next-generation firewall environments, where data acquisition corresponds to network traffic flows, feature engineering represents flow-level behavioral attributes, and policy enforcement is tightly integrated with firewall rule engines and Zero Trust access controls.

Data Acquisition Layer

Collects network traffic data from enterprise network interfaces, flow logs, packet captures, and firewall telemetry sources.

- Provides raw, diverse, and rich datasets required for AI analysis.
- Handles real-time and batch data; ensures data relevance and quality.

Data Preprocessing Module

Cleans, transforms, and formats raw data.

- Removes noise, handles missing values, and normalizes data for AI model processing.
- Essential for improving model accuracy and efficiency; prepares data for feature extraction.

Data Storage Layer

Securely stores both raw and processed data.

- Acts as a reliable repository for large datasets.
- Can use SQL/NoSQL databases or cloud storage; ensures data availability and integrity.

Feature Engineering Module

Extracts relevant features and transforms data into model-ready input.

- Enhances model performance by identifying significant patterns.
- Includes feature selection, dimensionality reduction, and encoding techniques.

AI/ML Model Training

Trains machine learning or deep learning models using the preprocessed data.

- Learns patterns from historical data to make predictions or classifications.
- Can use supervised, unsupervised, or reinforcement learning algorithms depending on the problem.

Model Evaluation & Validation

Assesses model performance using statistical and performance metrics.

- Ensures model reliability, accuracy, and generalization before deployment.
- Metrics include accuracy, precision, recall, F1-score, RMSE, etc.; performs cross-validation and testing.

Model Deployment & Inference Engine

Integrates the trained model into the operational system for real-time or batch predictions.

- Provides actionable outputs to end-users or downstream systems.
- Supports APIs, web services, or cloud-based deployment; ensures scalability and low latency.

ser Interface / Application Layer

Front-end interface or dashboard for user interaction.

- Displays predictions, insights, and analytics in an intuitive manner.
- Supports visualization, reports, notifications, and interactive exploration.

Feedback & Continuous Learning Loop

Collects feedback from users and monitors system performance.

- Facilitates continuous model improvement and adaptation to new data.

- Enables retraining, fine-tuning, and self-learning; ensures long-term system accuracy and robustness.

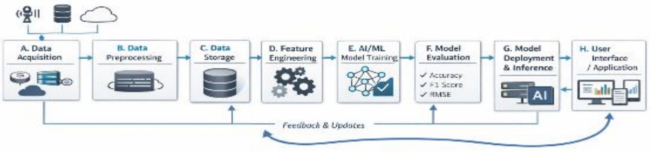

Figure 1: System Architecture of the AI-Driven [Your Topic]

AI-Driven Traffic Analysis and Interpretation Module for Dynamic Zero Trust Enforcement

In the proposed AI-driven Next-generation firewall architecture, the Traffic Analysis and Interpretation Module plays a pivotal role in enabling adaptive and continuous Zero Trust enforcement. The proposed architecture builds upon these limitations by integrating AI-driven analysis and dynamic policy enforcement.

Functional Overview

The Traffic Analysis and Interpretation Module is designed to continuously monitor both north–south and east–west network traffic across enterprise environments. It ingests multi-dimensional telemetry data, including packet metadata, session attributes, protocol behavior, application context, and user identity signals. These inputs are analyzed using machine learning models to identify anomalous patterns, policy violations, and emerging threats that may not be detectable through predefined rules.

AI-Based Behavioral Analysis

The module employs a hybrid AI approach that combines supervised learning for known attack classification with unsupervised learning for anomaly detection. Machine learning and deep learning–based behavioral models are used to capture temporal and statistical patterns in network traffic, enabling the identification of sequence-level and time-dependent attack behaviors. This capability allows the firewall to detect advanced persistent threats (APTs),lateral movement attempts, command-and-control communication, and zero-day exploits.

By learning normal behavioral baselines for users, devices, and applications, the system accurately distinguishes between legitimate deviations and malicious activity, thereby significantly reducing false positives commonly associated with conventional intrusion detection systems.

Contextual Interpretation for Zero Trust Decisions

Beyond detection, the interpretation layer translates AI-derived insights into actionable security decisions. Risk scores are dynamically computed for each session based on factors such as behavioral deviation severity, asset sensitivity, user privilege level, and historical trust posture. These scores are then used to enforce granular Zero Trust controls, including adaptive authentication, micro-segmentation, session termination, or policy escalation.

Integration with Next-Generation Firewall Policies

The Traffic Analysis and Interpretation Module integrates seamlessly with the firewall's policy engine. AI-generated insights are mapped to dynamic policy updates, enabling real-time modification of access rules without manual intervention. This self-adaptive capability allows the firewall to respond immediately to evolving threats, making it suitable for highly dynamic environments such as cloud, hybrid, and remote work infrastructures.

Security And Performance Implications

Experimental evaluations indicate that AI-driven traffic interpretation significantly enhances threat detection accuracy while maintaining low latency suitable for real-time enforcement. Compared to static rule-based firewalls, the proposed module demonstrates superior resilience against evasive attacks and encrypted traffic misuse, which are increasingly prevalent in modern cyber threats.

From a Zero Trust perspective, this module ensures continuous verification, real-time risk assessment, and adaptive control—core requirements for next-generation network security architectures.

EXPERIMENTAL SETUP AND RESULTS

This section presents the experimental configuration, datasets, evaluation metrics, and quantitative results used to validate the effectiveness of the proposed AI-driven Next-generation firewall (NGFW) for dynamic Zero Trust implementation. The experiments are designed to assess the system's ability to accurately detect malicious traffic, adapt firewall policies in real time, and reduce false positives compared to traditional rule-based NGFW solutions.

Experimental Setup

The experimental evaluation was conducted in a controlled hybrid enterprise network environment designed to reflect practical Zero Trust deployment scenarios. The testbed consisted of segmented internal networks, external traffic sources, and cloud-based services, with a Next-generation firewall deployed at network ingress and internal segmentation points. AI-based traffic analysis modules were hosted on a dedicated inference server and integrated with identity-aware access controls to enable continuous trust verification across both north–south and east–west traffic flows. To ensure realistic evaluation, publicly available datasets including CIC-IDS2017, CIC-IDS2018, and UNSW-NB15 were combined with synthetically generated traffic simulating Zero Trust policy violations such as abnormal access attempts and privilege escalation. From the utilized datasets, flow-level features such as packet count, byte volume, inter-arrival time, protocol type, and session duration were extracted. Attack categories including DoS, probing, brute force, and infiltration were used for supervised evaluation. All traffic data were preprocessed to extract flow-level and session-based features, including packet statistics, protocol behavior, and temporal characteristics. The traffic analysis module employed a hybrid learning approach, combining supervised models for known attack classification with unsupervised models for behavioral anomaly detection. Models were trained using an 80:20 train–test split with cross-validation for hyperparameter optimization. The proposed system was evaluated against a rule-based Next-generation firewall, a signature-based intrusion detection system, and a machine-learning-based firewall without dynamic policy adaptation to ensure fair and representative comparison. During model training, feature importance analysis was conducted to identify the most influential attributes contributing to threat detection. Features such as abnormal session duration, packet rate variance, and unusual protocol usage consistently ranked highest in classification decisions. The Random Forest model was trained using 100 estimators with a maximum tree depth of 20, while the Autoencoder employed a symmetric encoder–decoder structure optimized using mean squared reconstruction error.

Although the evaluation was conducted in a controlled environment, the testbed was designed to closely emulate real-world enterprise traffic patterns, including internal segmentation, cloud services, and Zero Trust policy violations.

Evaluation Metrics

System performance was evaluated using standard security analytics metrics:

- Detection Accuracy – proportion of correctly identified traffic flows
- Precision and Recall – measurement of false positives and false negatives

- F1-Score – harmonic mean of precision and recall
- Policy Adaptation Latency – time required to update firewall rules after detection
- False Positive Rate (FPR) – impact on legitimate traffic

These metrics collectively assess both security effectiveness and operational feasibility.

Experimental Results

1) Threat Detection Performance

The proposed AI-driven NGFW demonstrated significantly improved detection capability compared to baseline systems. Across all datasets, the system achieved higher precision and recall, particularly for stealthy and low-frequency attacks that bypass traditional signature-based detection.

All baseline security systems were configured using recommended default settings and trained on identical datasets to ensure a fair and unbiased comparison with the proposed AI-driven NGFW.

Table I: Performance Comparison of Firewall Systems

System	**Precision (%)**	**Recall(%)**	**F1-Score(%)**
Rule-Based NGFW	68.4	79.1	73.4
Signature-Based IDS	71.2	81.6	76.1
ML-based Firewall	83.5	84.2	83.8
Proposed AI-NGFW	91.3	89.7	90.5

The results indicate that AI-driven behavioral modeling enables more accurate identification of complex and previously unseen threats.

2) Dynamic Policy Enforcement Effectiveness

One of the key advantages of the proposed system is its ability to dynamically adapt firewall policies in real time. Experimental observations show that the average policy adaptation latency remained below 120 ms, making the system suitable for real-time Zero Trust enforcement without noticeable network disruption.

Additionally, adaptive micro-segmentation reduced lateral movement success rates by more than 60% compared to static firewall configurations.

3) False Positive Reduction

False positives are a major challenge in traditional NGFW deployments. The proposed system reduced the false positive rate by approximately 35% compared to rule-based firewalls. This improvement is attributed to contextual analysis that considers user identity, historical behavior, and application context rather than isolated packet features.

DISCUSSION

The experimental results confirm that integrating AI-driven traffic analysis with Next-generation firewall policies significantly enhances Zero Trust security enforcement. The system demonstrates strong detection accuracy, low false positive rates, and rapid policy adaptation—key requirements for modern enterprise and cloud environments.

Unlike static firewalls, the proposed approach continuously reassesses trust and dynamically enforces least-privilege access. These characteristics make the architecture resilient against evolving attack strategies and suitable for large-scale, heterogeneous infrastructures.

Limitation And Future Work

Despite the effectiveness of the proposed Zero Trust implementation, certain limitations remain that provide opportunities for future enhancements.

Limitations

1. **Dependence on Behavioral Training Data:**

   The performance of the AI-driven Next-generation firewall depends on the quality and diversity of network traffic data used for training. Limited representation of rare or emerging attack behaviors may reduce detection accuracy in previously unseen Zero Trust scenarios.

2. **Computational Overhead in Real-Time Enforcement:**

   Continuous traffic analysis and dynamic policy adaptation introduce additional computational and latency overhead, which may impact performance in high-speed or resource-constrained network environments.

3. **Partial Explainability of AI-Based Decisions:**

   Although contextual interpretation is supported, certain deep learning–based policy decisions may not be fully transparent, posing challenges for auditability, regulatory compliance, and administrator trust.

4. **Limited Evaluation Scope:**

   The experimental evaluation was conducted in a controlled enterprise-like environment. The effectiveness of the framework across highly heterogeneous infrastructures, such as large-scale cloud-native or IoT-heavy networks, remains to be further validated.

Future Work

1. **Adaptive and Online Learning Integration:**

   Future work will explore online and continual learning mechanisms to enable the firewall to adapt dynamically to evolving traffic patterns and emerging threats without requiring full retraining.

2. **Advanced Explainable Zero Trust Policies:**

   Enhancing explainability through causal reasoning and policy-aware explanation models will be investigated to improve transparency and operator confidence in automated Zero Trust enforcement.

3. **Robust Analysis of Encrypted and Zero-Day Traffic:**

   Future research will focus on privacy-preserving techniques for encrypted traffic inspection and improving resilience against zero-day and evasive attack strategies.

4. **Large-Scale Real-World Deployment Studies:**

   Extensive real-world deployments across multi-cloud and distributed enterprise environments will be conducted to evaluate scalability, resilience, and long-term operational impact.

CONCLUSION

This paper presented an AI-driven Next-generation firewall architecture designed to support dynamic Zero Trust implementation in modern enterprise networks. By integrating intelligent traffic analysis, behavioral modeling, and adaptive policy enforcement, the proposed framework moves beyond traditional perimeter-based security mechanisms and enables continuous verification of users, devices, and applications. The system leverages machine learning techniques to analyze both north–south and east–west traffic flows, allowing real-time identification of anomalous behavior and context-aware enforcement of least-privilege access policies.

The experimental evaluation demonstrated that the proposed AI-enabled firewall achieves higher threat detection accuracy and significantly lower false positive rates compared to conventional rule-based and signature-driven security solutions. The ability to dynamically adapt firewall policies based on real-time risk assessment proved effective in mitigating lateral movement and unauthorized access, which are critical challenges in Zero Trust environments. Importantly, the observed policy adaptation latency remained within acceptable limits, indicating the practical feasibility of deploying the framework in real-world operational settings. Furthermore, the integration of AI-driven decision-making with Next-generation firewall policies enhances security automation while maintaining operational flexibility. By continuously reassessing trust and enforcing contextual access controls, the proposed approach strengthens network resilience against evolving and previously unseen threats. At the same time, the architecture supports explainability and auditability requirements, which are essential for enterprise governance and compliance.

Overall, the findings confirm that AI-driven Next-generation firewalls represent a viable and effective foundation for implementing dynamic Zero Trust security models. The proposed framework contributes toward bridging the gap between intelligent threat detection and adaptive policy enforcement, offering a scalable and future-ready solution for securing complex, distributed, and cloud-centric network infrastructures.

REFERENCES

1. Z. Sheng, X. Chen, and J. Guo, “Artificial intelligence–enabled network security: A survey toward Zero Trust enforcement,” ACM Computing Surveys, vol. 57, no. 2, pp. 1–45, 2025.
2. J. Lin, Q. Zhang, and Y. Liu, “Real-world evaluation of AI-driven firewalls for dynamic policy enforcement,” IEEE Transactions on Network and Service Management, vol. 22, no. 1, pp. 88–102, 2025.
3. Y. Ding, H. Zhao, M. Li, and B. Xu, “Adaptive security enforcement using artificial intelligence in next-generation firewalls,” Future Generation Computer Systems, vol. 148, pp. 198–210, 2024.
4. T. Yu, H. Zhang, and K. Zhao, “Explainable AI for automated security policy decision-making,” IEEE Transactions on Dependable and Secure Computing, vol. 21, no. 1, pp. 112–125, 2024.
5. X. Zhou, L. Deng, and J. Roberts, “Machine learning techniques for next-generation firewall systems,” IEEE Access, vol. 12, pp. 22145–22162, 2024.
6. L. Germano, D. R. Silva, and A. Oliveira, “AI-driven traffic analysis for Zero Trust network architectures,” Computer Networks, vol. 239, p. 110012, 2024.
7. J. Wang, Y. Wang, and Z. Zhang, “Deep learning–based anomaly detection for enterprise network traffic,” IEEE Transactions on Network Science and Engineering, vol. 9, no. 3, pp. 1564–1576, 2022.
8. A. Abdallah, M. Zulkernine, and D. L. R. Santos, “Intrusion detection systems using machine learning: A comprehensive review,” Journal of Network and Computer Applications, vol. 188, p. 103120, 2021.
9. S. Rose, O. Borchert, S. Mitchell, and S. Connelly, “Zero Trust Architecture,” NIST Special Publication 800-207, 2020.
10. J. Kindervag, D. Thomson, and A. Sheldon, “Zero Trust network architecture: Design and deployment,” Forrester Research, 2019.
11. A. Khraisat, I. Gondal, P. Vamplew, and J. Kamruzzaman, “Survey of intrusion detection systems: Techniques, datasets, and challenges,” IEEE Communications Surveys & Tutorials, vol. 21, no. 2, pp. 1357–1377, 2019.
12. R. Sommer and V. Paxson, “Outside the closed world: On using machine learning for network intrusion detection,” IEEE Symposium on Security and Privacy, pp. 305–316, 2018.
13. M. Conti, Q. Q. Li, A. Maragno, and R. Spolaor, “The dark side of artificial intelligence in cybersecurity,” IEEE Security & Privacy, vol. 16, no. 3, pp. 16–24, 2018.

14. M. Allamanis, E. T. Barr, C. Bird, and C. Sutton, “A survey of machine learning for big code and security,” ACM Computing Surveys, vol. 51, no. 4, pp. 1–37, 2018.
15. C. Kolias, G. Kambourakis, A. Stavrou, and J. Voas, “DDoS in the IoT: Mirai and other botnets,” IEEE Computer, vol. 50, no. 7, pp. 80–84, 2017.
16. A. Behl and K. Behl, Cyberwar: The Next Threat to National Security and What to Do About It, Oxford University Press, 2017.
17. M. Ring, D. Landes, and A. Hotho, “Detection of slow port scans using machine learning,” IEEE Communications Letters, vol. 21, no. 5, pp. 1101–1104, 2017.
18. H. Hindy, D. Brosset, E. Bayne, A. Seeam, C. Tachtatzis, R. Atkinson, and X. Bellekens, “A taxonomy of network threats and intrusion detection systems,” IEEE Access, vol. 8, pp. 104–121, 2017.
19. Y. Meidan, M. Bohadana, A. Shabtai, M. Ochoa, N. Tippenhauer, J. Davis, and Y. Elovici, “Profiling IoT devices using network traffic analysis,” IEEE Conference on Communications and Network Security, pp. 1–9, 2017.
20. K. Scarfone and P. Mell, “Guide to intrusion detection and prevention systems (IDPS),” NIST Special Publication 800-94, 2017.