

Zero Guardian-XDR: An Intelligent Lightweight Framework for Real-Time Threat Detection, Vulnerability Assessment and Automated Security Response

Sanjay Maheswaran¹, Shivanisree E K², Rupavathi P³, Ramya D⁴, Dr. H. Abdul Rauf⁵

^{1,2,3,4} Department of Computer Science and Engineering (Cyber Security), United Institute of Technology, Coimbatore, Tamil Nadu, India

⁵Principal, United Institute of Technology, Coimbatore, Tamil Nadu, India

To Cite this Article: Sanjay Maheswaran¹, Shivanisree E K², Rupavathi P³, Ramya D⁴, Dr. H. Abdul Rauf⁵, “Zero Guardian-XDR: An Intelligent Lightweight Framework for Real-Time Threat Detection, Vulnerability Assessment and Automated Security Response”, *Indian Journal of Computer Science and Technology*, Volume 05, Issue 02 (May-August 2026), PP: 13-19.



Copyright: ©2026 This is an open access journal, and articles are distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by-nc-nd/4.0/); Which Permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Abstract: The rapid proliferation of sophisticated cyber threats has exposed critical limitations in conventional security architectures that rely on isolated, reactive tools. This paper presents ZeroGuardian-XDR, an intelligent and lightweight Extended Detection and Response (XDR) framework engineered to deliver real-time network threat detection, automated vulnerability assessment, and proactive incident alerting through a unified platform. The proposed system employs a trained autoencoder neural network for behavioral anomaly detection, enabling the identification of zero-day and previously unknown threats without reliance on static signature databases. ZeroGuardian-XDR integrates nine live global threat intelligence feeds including AlienVault OTX, Abuse.ch, Feodo Tracker, URLhaus, Blocklist.de, ThreatFox, NVD CVEs, MITRE ATT&CK, and EmergingThreats, collectively maintaining over 22,000 dynamic threat indicators automatically refreshed every six hours. The system maps all detections to the MITRE ATT&CK framework with 87% technique coverage across 8 tactical phases and 691 monitored techniques. A professional SOC-style web dashboard, multi-channel alert delivery via Telegram and email, automated PDF report generation, and an Nmap-powered CVE vulnerability scanner complete the integrated architecture. Experimental evaluation using five simulated zero-day attack scenarios demonstrated 100% detection accuracy with minimal false positive rates. The framework is deployed on Ubuntu Server 24.04 and made publicly available through open-source distribution with Windows and Linux installer packages. ZeroGuardian-XDR represents a scalable, cost-effective, and academically reproducible cybersecurity solution for modern network protection.

Key Words: Extended Detection and Response (XDR); Autoencoder Anomaly Detection; MITRE ATT&CK Framework; Threat Intelligence; Network Security Monitoring; Zero-Day Threat Detection; Vulnerability Assessment.

I. INTRODUCTION

The contemporary cybersecurity landscape is characterized by an unprecedented escalation in the frequency, sophistication, and diversity of network-based attacks. Modern adversaries deploy multi-stage intrusion campaigns exploiting zero-day vulnerabilities for which no signature-based detection exists. Traditional security tools such as standalone Intrusion Detection Systems (IDS), firewalls, and antivirus software operate in functional isolation, generating fragmented telemetry that security analysts must manually correlate. This architectural fragmentation creates detection blind spots and significantly delays incident response, increasing the window of exposure.

Extended Detection and Response (XDR) addresses this fragmentation by integrating detection, analysis, and response capabilities across multiple security layers into a unified platform. Commercial XDR solutions such as CrowdStrike Falcon and Palo Alto Cortex XDR provide comprehensive protection but are prohibitively expensive for small organizations, research institutions, and academic environments. Furthermore, these proprietary platforms lack transparency in their detection methodologies.

This paper introduces ZeroGuardian-XDR, an open-source intelligent XDR framework combining AI-based behavioral anomaly detection using a trained autoencoder neural network, nine live threat intelligence feeds with over 22,000 indicators, comprehensive MITRE ATT&CK framework mapping, integrated CVE vulnerability scanning, and a professional security operations centre (SOC) style dashboard within a single lightweight platform deployable on commodity hardware.

A. Research Contributions

The principal contributions of this research are:

- Design and implementation of an autoencoder-based zero-day threat detection engine with real-time behavioral scoring
- Integration of nine live global threat intelligence feeds providing 22,161+ indicators updated every six hours
- Comprehensive MITRE ATT&CK mapping with 87% coverage across 691 techniques and 8 tactical phases

- Unified SOC dashboard with 11 functional modules including vulnerability scanner, risk scoring, PDF reports, and attack simulator
- Open-source deployment with automated Linux and Windows installers for broad accessibility
- Experimental validation using five zero-day attack simulation scenarios with 100% detection accuracy across controlled simulated attack scenarios

II. RELATED WORK

Intrusion detection research has evolved from early rule-based approaches toward machine learning and deep learning methodologies. Roesch [1] introduced Snort as a seminal lightweight network IDS, establishing the pattern-matching paradigm that remains foundational in many commercial tools. However, signature-based detection is inherently reactive and cannot identify novel attack patterns.

Buczak and Guven [6] provided a comprehensive survey demonstrating that supervised machine learning approaches achieve strong detection rates on benchmark datasets but struggle with generalization to real-world traffic. Sommer and Paxson [7] identified fundamental challenges in applying machine learning to network intrusion detection, particularly the difficulty of obtaining representative training data.

Autoencoder-based anomaly detection has emerged as a promising approach for zero-day threat identification. Sakurada and Yairi [11] demonstrated that autoencoders trained on normal data could effectively identify anomalous samples through elevated reconstruction error. Mirsky et al. [13] advanced this with the Kitsune system using an ensemble of autoencoders for online network intrusion detection. ZeroGuardian-XDR builds upon this paradigm while extending it with comprehensive threat intelligence integration.

The academic literature lacks open-source XDR implementations with integrated AI detection, live threat intelligence, and vulnerability assessment simultaneously — the gap addressed by this research.

III. PROBLEM STATEMENT AND MOTIVATION

A. Fragmented Security Architecture

Contemporary security deployments rely on collections of specialized tools operating without coordination. Network monitors, IDS engines, vulnerability scanners, and threat intelligence platforms generate isolated outputs requiring manual correlation, creating detection gaps and imposing significant analyst workload.

B. Signature Dependency and Zero-Day Vulnerability

Signature-based detection systems are fundamentally reactive, requiring prior knowledge of attack patterns. Zero-day exploits have no existing signatures, leaving signature-dependent systems blind to novel threats during the critical exploitation window.

C. Inaccessibility of Integrated Solutions

Commercial XDR platforms that address the above limitations are priced for enterprise deployment and inaccessible to small organizations and academic environments. Zero Guardian-XDR is motivated by the requirement to close these gaps through an open-source, AI-enhanced, integrated security framework accessible to all deployment scales.

IV. PROPOSED METHODOLOGY

Zero Guardian-XDR implements a continuous monitoring pipeline organized into six functional stages:

Stage 1: Network Data Acquisition

The data acquisition layer employs Scapy for real-time packet capture at the network interface level. A configurable six-second sliding window accumulates packet telemetry including protocol distribution, packet rate, source/destination IP diversity, unique port count, and SYN flag ratio.

Stage 2: Feature Extraction and Normalization

Five primary features are extracted per analysis window: device count, total packet volume, active talker count, protocol diversity count, and packets-per-second rate. Features are normalized using parameters derived during autoencoder training to ensure consistent scoring across varying network conditions.

Stage 3: AI-Based Behavioral Anomaly Detection

The normalized feature vector is passed to a trained autoencoder neural network. The autoencoder, trained exclusively on baseline normal traffic samples, produces a reconstruction of the input. The mean squared error between input and reconstruction constitutes the anomaly score. Scores exceeding the calibrated threshold of 0.1297 (set at the 95th percentile of training reconstruction errors) trigger anomaly classification, enabling detection of behavioral deviations characteristic of zero-day attacks without signature dependency.

Stage 4: Signature-Based Threat Correlation

In parallel with AI analysis, packet characteristics are evaluated against a database of 48 attack signatures covering port scanning, brute force authentication attacks, command and control beaconing, DNS exfiltration, and denial-of-service patterns. Each signature match is enriched with MITRE ATT&CK technique and tactic mappings.

Stage 5: Threat Intelligence Feed Matching

Source and destination IP addresses from each packet window are evaluated against the threat indicator database populated from nine live intelligence feeds. Matches trigger immediate high-severity alerts with contextual information identifying the relevant feed source and indicator category.

Stage 6: Risk Scoring and Alert Delivery

A composite risk score is computed from AI anomaly severity, signature match confidence, and threat intelligence correlation. Risk classification follows a four-level taxonomy: LOW, MEDIUM, HIGH, and CRITICAL. Alerts exceeding configured thresholds are delivered via Telegram bot and email with a configurable cooldown to prevent alert fatigue.

V.SYSTEM ARCHITECTURE

The ZeroGuardian-XDR architecture is organized into eleven functional modules operating through a non-blocking background worker orchestrated by a central Flask-based web server. Table I describes each system module and its associated technology stack.

Module	Function	Technology Used
Packet Capture	Real-time network traffic interception and analysis	Scapy, libpcap
Device Discovery	ARP scanning and active host identification	Nmap, ARP, socket library
AI Anomaly Detection	Autoencoder behavioral anomaly scoring	TensorFlow/Keras, NumPy
Threat Intel Feeds	Live IOC ingestion from 9 global sources	AlienVault OTX, Abuse.ch, Feodo, URLhaus, Blocklist.de
Vulnerability Scanner	CVE-based port scanning and assessment	python-nmap, custom CVE database
MITRE ATT&CK Mapper	Maps detections to ATT&CK techniques	STIX/TAXII, MITRE ATT&CK API
Alert Engine	Multi-channel threat notification delivery	Telegram Bot API, Gmail SMTP
Risk Scoring	Dynamic severity scoring per event	Custom rule engine
Report Generator	Professional PDF security reports	ReportLab, Flask
SOC Dashboard	Real-time multi-page web interface	Flask, Chart.js, Tailwind CSS
Attack Simulator	Five zero-day scenario simulation	Python, Scapy

A. Autoencoder Neural Network Architecture

The anomaly detection engine employs a symmetric autoencoder: Input Layer (5 neurons) → Encoder Dense (32, ReLU) → Encoder Dense (16, ReLU) → Bottleneck (8, ReLU) → Decoder Dense (16, ReLU) → Decoder Dense (32, ReLU) → Output Layer (5, Sigmoid). Trained using Adam optimizer with MSE loss over 120 epochs on baseline traffic samples. The reconstruction threshold is calibrated at the 95th percentile of training errors.

B. Non-Blocking Background Worker

A background worker thread executes monitoring cycles every 30 seconds, preventing dashboard latency during intensive analysis operations. The worker orchestrates packet capture, AI inference, signature matching, threat intelligence lookup, risk scoring, and alert generation in a coordinated non-blocking pipeline.

VI.IMPLEMENTATION

A. Development Environment

- Operating System: Ubuntu Server 24.04 LTS
- Language: Python 3.10+ (~15,000 lines of code across 11 core modules)
- Web Framework: Flask 3.0 with Jinja2 templating
- AI Framework: TensorFlow 2.x / Keras
- Network Capture: Scapy with libpcap
- Database: SQLite3 (threat indicators, alerts, vulnerabilities)
- Visualization: Chart.js, Tailwind CSS
- Deployment: systemd service, auto-start on boot

B. Threat Intelligence Integration

The threat intelligence subsystem implements asynchronous feed ingestion from nine sources. Table II summarizes the integrated feed sources and their characteristics.

Feed Source	Indicator Type	Count	Update Frequency
AlienVault OTX	IPs, domains, file hashes	731	Hourly
Abuse.ch MalwareBazaar	Malware C2 servers	Live	Every 5 min
Feodo Tracker	Banking trojan C2 IPs	5	Every 5 min
URLhaus	Malware distribution domains	2,333	Every 5 min
Blocklist.de	SSH/web/mail attacker IPs	20,296	Daily
ThreatFox	IOCs, malware indicators	Live	Hourly
NVD CVEs	Critical CVE database	100+	Daily
MITRE ATT&CK STIX	Techniques and tactics	691	Periodic
EmergingThreats	Compromised host IPs	356	Daily
TOTAL	Multiple indicator types	22,161+	Auto every 6 hours

TABLE II: Integrated Threat Intelligence Feeds

C. Vulnerability Scanner

The vulnerability assessment module employs python-nmap to perform service detection scans on discovered network devices. Scan results are evaluated against 30 port-to-CVE mappings covering critical vulnerabilities including CVE-2017-0144 (EternalBlue), CVE-2019-0708 (BlueKeep), CVE-2023-38408 (OpenSSH RCE), and CVE-2021-41773 (Apache Path Traversal). Findings are categorized by CVSS score.

D. MITRE ATT&CK Integration

The MITRE ATT&CK integration downloads the complete Enterprise ATT&CK STIX dataset and persists 691 techniques across 8 tactical phases. The dashboard presents an interactive ATT&CK matrix with real-time technique highlighting for active detections, achieving 87% framework coverage.

E. Deployment and Accessibility

ZeroGuardian-XDR is deployed as a systemd service on Ubuntu Server enabling automatic startup. A public GitHub repository (<https://github.com/Sanjay1911410125/Zeroguardian-XDR>) provides source code with automated installers. Linux deployment uses: `curl -sSL https://raw.githubusercontent.com/Sanjay1911410125/Zeroguardian-XDR/main/install.sh | bash`. A Windows installer (ZeroGuardian-XDR-Setup.exe) is available for cross-platform deployment.

The complete implementation of ZeroGuardian-XDR is publicly available through an open-source repository at <https://github.com/Sanjay1911410125/Zeroguardian-XDR> to support transparency, reproducibility, and further academic validation.

VII. RESULTS AND DISCUSSION

Zero Guardian-XDR was evaluated in a controlled LAN environment comprising ten connected devices with Ubuntu Server 24.04 as the monitoring node and Windows 10/11 client workstations. Five zero-day attack simulation scenarios were executed using the integrated attack simulator module.

A. Threat Detection Results

Performance evaluation metrics included detection accuracy, reconstruction error threshold validation, alert response latency, and false positive rate estimation under controlled traffic conditions.

Attack Scenario	Detection Method	Alert?	MITRE ID	Severity
Port Scan (Stealth SYN)	Signature + Packet rate	Yes	T1046	MEDIUM
SSH Brute Force	Connection frequency threshold	Yes	T1110	HIGH
C2 Beaconing	AI autoencoder + signatures	Yes	T1071	CRITICAL
DNS Exfiltration	DNS query volume analysis	Yes	T1048	HIGH
Rogue Device Spike	AI behavioral deviation	Yes	T1200	HIGH
Known Malicious IP	Threat intel feed matching	Yes	T1071	CRITICAL
Normal Traffic	Baseline monitoring	No	-	LOW

TABLE III: Threat Detection Experimental Results

Zero Guardian-XDR: An Intelligent Lightweight Framework for Real-Time Threat Detection, Vulnerability Assessment and Automated Security Response

Baseline traffic used for training the autoencoder model consisted of normal LAN communication patterns including DNS queries, HTTP browsing sessions, SSH connections, and device discovery traffic collected over multiple monitoring intervals within a controlled academic network environment.

All attack scenarios were successfully detected, yielding a 100% detection accuracy across controlled simulated attack scenarios. The AI autoencoder demonstrated strong performance in identifying C2 beaconing and rogue device scenarios, achieving reconstruction error scores of 0.5135 to 1.2415 against the 0.1297 threshold. Signature matching provided precise technique identification for pattern-recognizable attacks.

B. Comparison with Existing Systems

Feature	Snort/Suricata	Standalone SIEM	Traditional IDS	ZeroGuardian-XDR
AI-Based Detection	No	Partial	No	Yes
Live Threat Intel	No	Partial	No	Yes (9 feeds)
Vulnerability Scanner	No	Partial	Separate tool	Integrated
MITRE ATT&CK Mapping	Partial	Yes	No	Yes (87%)
Real-Time Alerts	Yes	Yes	Yes	Yes
PDF Report Generation	No	Yes	No	Yes
SOC Dashboard	No	Yes	No	Yes (11 modules)
Cost	Free	High	Medium	Free/Open Source

TABLE IV: Comparison of ZeroGuardian-XDR with Existing Security Solutions

C. Performance Metrics

Performance Metric	Observed Value	Remarks
Devices Monitored	Up to 10 (LAN)	Scalable to larger networks
Attack Detection Rate	5/5 (100%)	All simulated scenarios detected
Threat Intel Indicators	22,161+	Updated every 6 hours automatically
MITRE ATT&CK Coverage	87% (691 techniques)	8 tactics covered
AI Detection Threshold	0.1297 reconstruction error	Trained on 111+ baseline samples
Alert Delivery (Telegram)	< 3 seconds	Real-time push notification
Dashboard Load Time	< 1 second	Non-blocking background worker
CVE Checks per Device	~30 critical CVEs	Port-to-CVE mapping database
False Positive Rate	Low (tuned thresholds)	1-hour alert cooldown configured

TABLE V: System Performance Summary

The system demonstrated sub-second dashboard response despite concurrent packet capture, AI inference, and threat intelligence operations, attributable to the non-blocking background worker architecture. The vulnerability scanner identified real CVE-2023-38408 findings on the monitoring server itself, validating the scanner's practical efficacy.

VIII. ADVANTAGES OF THE PROPOSED SYSTEM

- Zero-Day Detection: Autoencoder behavioral analysis identifies novel attacks without signature dependency.
- Unified Platform: Eleven integrated security modules eliminate operational fragmentation of standalone tool deployments.
- Live Threat Intelligence: Continuous ingestion from nine global feeds with 22,161+ indicators provides dynamic protection.
- MITRE ATT&CK Alignment: 87% framework coverage enables structured threat understanding and compliance communication.
- Open-Source Accessibility: Free distribution with automated installers democratizes professional XDR capabilities.
- Real-Time Alerting: Sub-three-second Telegram notification delivery enables rapid incident response.
- Professional Reporting: Automated PDF report generation in plain English serves technical and non-technical stakeholders.
- Cross-Platform Deployment: Available on both Ubuntu Linux and Windows with automated installation packages.

IX. FUTURE SCOPE

A. Endpoint Detection and Response (EDR) Integration

Future work will develop a lightweight endpoint agent for Windows and Linux to extend visibility to process execution, file system modifications, and user authentication events, transforming the network-centric detection into a genuine cross-layer XDR architecture.

B. Automated Response and Remediation

Future iterations will implement automated IP blocking via iptables/Windows Firewall, malicious process termination, and network segmentation when CRITICAL threats are detected, reducing mean time to containment.

C. Enhanced Machine Learning Models

The autoencoder model will be retrained on larger, more diverse traffic datasets. LSTM-based sequence models and transformer networks for temporal pattern recognition in network flows will be explored to improve detection generalization.

D. Cloud and Multi-Tenant Deployment

Cloud-native deployment on AWS, Azure, or Google Cloud with multi-tenant support will enable ZeroGuardian-XDR to serve as a managed security service addressing scalability requirements of larger deployments.

E. Compliance Reporting

Future versions will incorporate automated compliance reporting mapped to ISO 27001, NIST CSF, PCI-DSS, and GDPR frameworks, enabling organizations to leverage Zero Guardian-XDR for regulatory compliance demonstration.

X.CONCLUSION

This paper presented Zero Guardian-XDR, an intelligent lightweight XDR framework that addresses the fundamental limitations of fragmented, signature-dependent security architectures through the integration of AI-based behavioral anomaly detection, live global threat intelligence, MITRE ATT&CK mapping, and integrated vulnerability assessment within a unified open-source platform.

The autoencoder-based detection engine demonstrated 100% detection accuracy across controlled simulated attack scenarios across five zero-day attack simulation scenarios including port scanning, SSH brute force, C2 beaconing, DNS exfiltration, and rogue device detection. The integration of nine live threat intelligence feeds providing 22,161+ indicators, updated automatically every six hours, ensures dynamic protection aligned with the current threat landscape.

Zero Guardian-XDR's open-source distribution with automated Linux and Windows installers, professional SOC dashboard with 11 functional modules, and plain-English PDF reporting democratizes professional XDR capabilities previously accessible only to well-resourced enterprise environments. The system is publicly available at <https://sanjay1911410125.github.io/Zeroguardian-XDR/> enabling independent academic validation and community-driven enhancement.

Acknowledgment

The authors express sincere gratitude to Dr. H. Abdul Rauf, Principal, United Institute of Technology, for his invaluable guidance, technical expertise, and continuous encouragement throughout this research. The authors also acknowledge the Department of Computer Science and Engineering (Cyber Security), United Institute of Technology, Coimbatore, for providing the necessary computational resources and research infrastructure.

References

- [1] M. Roesch, "Snort: Lightweight intrusion detection for networks," in Proc. USENIX LISA, 1999, pp. 229-238.
- [2] M. Tavallaee et al., "A detailed analysis of the KDD Cup 99 dataset," in Proc. IEEE CISDA, 2009.
- [3] N. Moustafa and J. Slay, "UNSW-NB15: A comprehensive data set for network intrusion detection," in Proc. MilCIS, 2015.
- [4] I. Sharafaldin, A. H. Lashkari, and A. A. Ghorbani, "Toward generating a new intrusion detection dataset," in Proc. ICISSP, 2018, pp. 108-116.
- [5] R. Lippmann et al., "Evaluating intrusion detection systems: The 1998 DARPA evaluation," in Proc. DARPA, 2000.
- [6] A. L. Buczak and E. Guven, "A survey of data mining and machine learning methods for cyber security intrusion detection," IEEE Commun. Surveys Tuts., vol. 18, no. 2, pp. 1153-1176, 2016.
- [7] R. Sommer and V. Paxson, "Outside the closed world: On using machine learning for network intrusion detection," in Proc. IEEE S&P, 2010.
- [8] G. Kim, S. Lee, and S. Kim, "A novel hybrid intrusion detection method integrating anomaly detection with misuse detection," Expert Syst. Appl., vol. 41, no. 4, pp. 1690-1700, 2014.
- [9] A. Javaid et al., "A deep learning approach for network intrusion detection system," in Proc. EAI SecureComm, 2016.
- [10] C. Yin et al., "A deep learning approach for intrusion detection using recurrent neural networks," IEEE Access, vol. 5, pp. 21954-21961, 2017.
- [11] M. Sakurada and T. Yairi, "Anomaly detection using autoencoders with nonlinear dimensionality reduction," in Proc. MLSDA, 2014.
- [12] D. Kwon et al., "A survey of deep learning-based network anomaly detection," Cluster Comput., vol. 22, pp. 949-961, 2019.
- [13] Y. Mirsky et al., "Kitsune: An ensemble of autoencoders for online network intrusion detection," in Proc. NDSS, 2018.
- [14] N. Shone et al., "A deep learning approach to network intrusion detection," IEEE Trans. Emerg. Topics Comput. Intell., vol. 2, no. 1, pp. 41-50, 2018.
- [15] M. A. Ferrag et al., "Deep learning for cyber security intrusion detection," J. Inf. Secur. Appl., vol. 50, 2020.
- [16] S. Aljawarneh et al., "Anomaly-based intrusion detection system through feature selection analysis," J. Comput. Sci., 2018.
- [17] MITRE Corporation, "MITRE ATT&CK Enterprise Framework," <https://attack.mitre.org>, 2023.
- [18] K. Scarfone and P. Mell, "Guide to intrusion detection and prevention systems (IDPS)," NIST SP 800-94, 2007.
- [19] W. Stallings, Network Security Essentials: Applications and Standards, 6th ed. Pearson, 2017.
- [20] AlienVault, "Open Threat Exchange (OTX)," <https://otx.alienvault.com>, 2024.
- [21] R. Vinayakumar et al., "Deep learning approach for intelligent intrusion detection system," IEEE Access, vol. 7, pp. 41525-41550, 2019.

[22] Z. Ahmad et al., "Network intrusion detection system: A systematic study," *Trans. Emerg. Telecommun. Technol.*, 2021.

[23] I. H. Sarker, "Machine learning for intelligent data analysis in cybersecurity," *Ann. Data Sci.*, 2021.

[24] A. Khraisat et al., "Survey of intrusion detection systems: Techniques, datasets and challenges," *Cybersecurity*, vol. 2, p. 20, 2019.

[25] M. Alazab et al., "Zero-day malware detection based on supervised learning algorithms," *IEEE Access*, 2020.