FMDB *Transactions on Sustainable Intelligent Networks*

Optimizing Network Traffic Through AI-Enhanced Zero-Trust Architectures

S. Vishal[1,*], S. Sai Vishaal[2]

[1,2]Department of Artificial Intelligence and Machine Learning, SRM Institute of Science and Technology, Ramapuram, Chennai, Tamil Nadu, India.
vishalshan04@gmail.com[1], saivishaal2003@gmail.com[2]

Abstract: Zero-Trust Architectures (ZTA) migration involves transitioning from perimeter-based security to persistent authentication for every request to access resources. Although it introduces security, the migration comes with a high overhead cost, inducing network latency and increased complexity in policy administration. This paper presents an AI-based framework for ZTA that minimizes network traffic flow without compromising security tenets. Our solution involves a deep reinforcement learning (DRL) agent that dynamically varies network paths and access privileges in real-time depending on device posture, end-user behaviour, and app sensitivity. The study was proven in an emulated enterprise network deployment. The primary data used here is the 'ZTA-Traffic-Sim-2025', which contains simulated data of 10 million network flow records. Traffic for 5,000 users and 15,000 devices was simulated for a month, incorporating various simulated attack vectors. The model is trained using Python with TensorFlow for the DRL agent and the ns-3 simulator to simulate the network environment. The results confirm that the AI-based ZTA reduces average network latency by up to 35% and throughput by 25% compared to static ZTA, while improving the detection rate for anomaly activity by 18%.

Keywords: Zero-Trust Architecture; Artificial Intelligence; Network Traffic Optimization; Deep Reinforcement Learning; Device Posture; End-User Behaviour; App Sensitivity; Network Environment.

Received on: 03/09/2024, **Revised on:** 15/11/2024, **Accepted on:** 17/12/2024, **Published on:** 03/06/2025

Journal Homepage: https://www.fmdbpub.com/user/journals/details/FTSIN

DOI: https://doi.org/10.69888/FTSIN.2025.000380

Cite as: S. Vishal and S. S. Vishaal, "Optimizing Network Traffic Through AI-Enhanced Zero-Trust Architectures," *FMDB Transactions on Sustainable Intelligent Networks*., vol. 2, no. 2, pp. 59–68, 2025.

1. Introduction

Ali et al. [1] argued that the rate of digital innovation has accelerated due to the ubiquitous physical-layer deployment of mobile devices, IoT devices, and cloud computing, which has dissolved traditional network perimeters. The physical and conceptual castle-and-moat approach that worked well to secure static enterprise landscapes is no longer tenable today to secure highly dynamic, hybrid landscapes of contemporary organizations. These traditional architectures inherently depend on insiders, making them vulnerable to malicious insider credential misuse and lateral movement. With resources more distributed across clouds and mobile endpoints today, network location-based access to resources is insufficient and insecure. This led to a paradigm shift, which in turn constructed the cradle for Zero-Trust Architecture (ZTA), an ideology centered on the assumption

[*]Corresponding author.

that no one—internal or external—should be trusted by default. Annabi et al. [2] discussed the emergence of ZTA as a reaction to changing threat vectors and eroding trust in perimeter architectures.

The "never trust, always verify" principle is the foundation of ZTA, where enforcement and verification are persistent, regardless of the origin of a request. This step eliminates blind trust by enforcing user identity authentication, device health checks, location verification, and security posture checks on every access attempt. In contrast to traditional architecture, where authenticated logon grants broad access, ZTA imposes micro-segmentation on applications to limit lateral movement. To achieve this, transactional and context-based access control are required. The integration of controls during deployment time enables a secure yet adaptable infrastructure that is compatible with the distributed nature of modern enterprise systems. Asensio-Garriga et al. [3] developed an entirely automated security service level agreement lifecycle model for 5G networks, thus demonstrating how the principles of ZTA are applied in high-speed and decentralized environments. Their research demonstrates that real-time, enforced fine-grained access policies ensure that only approved, authenticated, and policy-conformant entities can communicate with sensitive resources. These policies are enforced dynamically, with security settings flexible enough to evolve in line with changing risk landscapes.

The key to this is the separation of location and identity; identity is the new border. This is a simultaneity-based methodology among ZTA-need identity providers, device management platforms, and threat detection engines. Machine learning approaches have been utilized by Awan et al. [4] to detect anomalous access behavior and mandate ZTA deployment on a baselining behavior schedule. From endpoint, server, and authentication mechanism logs, AI solutions can detect anomalies from typical user behavior that indicate compromise. The intelligent system makes ZTA models move beyond rule sets and become adaptive security architectures instead. Behavioral analytics offers access based on credentials, but also considers whether the access request aligns with established usage patterns. Active defense through this mechanism significantly raises the bar for insider and lateral attackers, who often attempt to disguise themselves as legitimate traffic. Bello et al. [5] also pointed to the practical costs of deploying ubiquitous ZTA. They demonstrated that continuous identity authentication, device conformity verification, and context indications in attempted access can cause latency and degrade the user experience. Security is improved, but at the expense of causing workflow disruptions and network performance degradation due to the introduced overhead. Furthermore, highly fine-grained policy enforcement in elastic environments imposes higher administrative overhead.

Each new user, application, or device interaction may invite a policy refresh. Without automation, this presents an open invitation to the risk of policy drift and misconfiguration, compromising the effectiveness of ZTA. Coronado et al. [6] proposed an AI-based orchestration strategy to solve ZTA's operational challenges. It employed network orchestration models and AI to dynamically refresh real-time threat categorizations and make responsive policy adjustments. Through real-time monitoring of telemetry from multiple sources—viz, access logs, firewalls, and intrusion detection systems—the AI engine can predict security incidents ahead of time. With the capability to predict, organizations can maintain high security levels without significant performance trade-offs. Orchestration layers also make policy enforcement simpler in hybrid and multi-cloud environments. Da Silva and Santos [7] elaborated on the complementarity of ZTA and smart authentication mechanisms. They had federated identity infrastructures and adaptive, risk-based authentication, which enabled the development of a real-time trust rating that could be used to make informed decisions about access. These trust ratings are based on measures like user activity, device reputation, and geolocation history. Where the user deviates from normal behavior—e.g., trying to log on from overseas—the system can demand further authentication or even outright refuse access.

This Layered, Adaptive complements ZTA's thrust towards eliminating implicit trust and enhancing contextual awareness wherever devices are used. Liu et al. [8] conducted early work on anomaly detection using unsupervised machine learning models, which remain the foundation of ZTA deployments to date. They had their models learn against monolithic telemetry data sets to establish baselines of the machines' and end users' behavioral patterns. They reported anomalies at the sites where they were identified as potential intrusions. These observations have been utilized to develop automated threat detection engines integrated into current commercial ZTA products. Live flagging and remediation are employed to reduce dwell time and subsequently decrease the chance of damage caused by insider and external attacks.

Sirohi et al. [9] used reinforcement learning models to generate dynamic enterprise network access policy decisions. The models are part of the chain of optimal decision-making that occurs after experiencing the environment and understanding the impact of the access control decision. It is a self-optimizing model that is adaptive in the adaptive ZTA environment, becoming stronger and more intelligent over time. Feedback loops regularly enforce posture to secure against new emerging threats that are discovered.

Patil et al. [10] further developed this concept by integrating trust scoring and policy recommendation engines to enhance decision-making in ZTA systems. Their hybrid engine relies on past access, real-time threat intelligence, and context to make highly accurate access decisions. These AI-powered ZTA systems enforce fine-grained policies without undue delay, providing

the optimal balance of security and usability. Through trustworthiness filtering, organizations can promote good behavior more effectively and investigate potential bad behavior more assertively.

Sengupta and Anantharaman [11] demonstrated that, through the combination of AI and ZTA, false positives for threat detection can be minimized and context awareness can be increased. Rule-based traditional security products generate a high volume of false alarms, leading to a backlog in the security operations center. AI would then distinguish between benign anomalies and true threats by taking into account the broader context of trends in behavior, thereby enhancing the enforcement effectiveness of ZTA. Contextual intelligence is also instrumental in avoiding alert fatigue and facilitating rapid incident response. Srinivas et al. [12] stated the need for the availability and integrity of data to facilitate successful ZTA deployments. AI-based solutions require enormous amounts of clean, up-to-date data upon which decisions are made. Network, server, and endpoint device telemetry must be available to report on the efficacy of policy enforcement automation and threat detection. They proved the need for scalable data gathering mechanisms to facilitate these functions.

DeCusatis et al. [13] established that ZTA frameworks must be interoperable and modular, allowing them to evolve in response to open policy requirements and emerging technologies. They facilitate the seamless integration of AI models, identity providers, and policy engines, yielding a unified and secure security posture. This prevents vendor lock-in and allows more flexibility in addressing dynamic compliance requirements or evolving threat models. Assunção [14] introduced a decentralized policy assessment framework that supports edge-of-network decision-making. From the decision that is most proximal to the trusted site, ZTA can actually minimize latency and decrease dependence on centralized sites. This type of framework is particularly common in IoT and edge computing scenarios, where real-time responsiveness is crucial. Remediation is also automated, utilizing the model, which enables the automatic recovery of systems. Lukaseder et al. [15] envisioned an extensive examination of AI-driven ZTA systems in mission-critical infrastructure scenarios. Their results supported that the technologies are not making notable contributions to mean time to detect (MTTD) and mean time to respond (MTTR) to attacks. They also demonstrated that risk analysis and behavior modeling based on automation can provide higher operating resilience without compromising regulatory compliance. Their study makes AI-driven ZTA a pillar of enterprise cybersecurity planning in the digital age.

2. Review of Literature

Annabi et al. [2] formulated a comprehensive understanding of the fundamentals of Zero-Trust Architecture (ZTA), with an emphasis on what distinguishes it from traditional perimeter-based methods. Their article categorizes the way in which the ZTA model redefines trust, not by location in the network but as being dynamic and constantly asserted. The architecture is based on three key components: secure access to resources regardless of a device's location within the network, least-privilege access enforcement, and continuous monitoring and auditing. Micro-segmentation is the norm in deployments, with the network segmented into hundreds of small areas known as segments in an attempt to restrict the lateral movement of an intrusion. Segments are managed by segmentation gateways, with greater control over traffic and halting lateral movement. The policy node, or decision-maker, contextually denies or permits access based on multiple context indicators, such as identity, device, and intent.

Asensio-Garriga et al. [3] proposed autonomous service-level control in next-generation networks, which aligns with the demand for Zero Trust's dynamic, contextual decision-making. With ZTA maturity, its deployment demands scalability and automation that conventional systems cannot achieve. They encompass architectural pattern designs that enable the enforcement and orchestration of policy systems to react in real-time, thus providing an instantaneous response to malicious activity. These frameworks constitute the prominent interstitial construct between theoretical ZTA notions and their implementation in state-of-the-art current environments, such as 5G networks and cloud-native systems. Awan et al. [4] utilized artificial intelligence to enhance detection against sophisticated cyberattacks, echoing the increased use of AI in Zero Trust environments. A machine learning application for dynamic behavioral inspection and threat discovery was also among their techniques. Anomaly detection, focusing on predetermined areas, utilizes unsupervised learning-enabled systems to automatically detect deviations from normal behavior. It is analogous to Zero Trust's ongoing authentication and offers an automatic response, i.e., quarantining a random device exhibiting suspicious behavior without affecting the entire network.

The authors of Bello et al. [5] employed AI and a deep learning model to support decision-making within cybersecurity systems. They utilized supervised models, which utilize tagged sets of malware and normal behaviors to accurately classify incoming streams of data. With such models, policy engines become wiser and smarter, with more informed decision logic over time. With such adaptability, proactive threat mitigations and strong policy feedback loops are enabled within Zero Trust systems. Early efforts in behavior-based anomaly detection have been provided by Liu et al. [8], who are also contributors to the core AI-reliant approaches to Zero Trust models. Their frameworks would be able to define baseline normal user and system standards, serving as markers for identifying zero-day vulnerabilities. The ability of the frameworks to execute with minimal,

or even without, continuous human input is a direct facilitator of ZTA's vision to reduce manual configuration and intervention, thereby providing dynamic and scalable security postures across distributed networks.

Sengupta and Anantharaman [11] had previously experimentally confirmed live AI deployment for adaptive authentication systems and had reported empirical results on dynamic access policy tuning. The authors' contributions support the use of context, such as geolocation, time, and device health, and dynamically adjusting the level of authentication per access request. It is this sort of control that delivers a high-grain impact on Zero Trust, which assumes compromise and dynamically verifies access through continuous risk verification instead of static credentials. Srinivas et al. [12] detailed the constraints of implementing Zero Trust into existing enterprise infrastructures, noting that the distributed architecture and legacy infrastructure created integration problems.

The study remains relevant today as companies strive to integrate ZTA and AI platforms into their evolving infrastructures. The research preferred modularity within architectures and necessitated compatibility layers that enable secure interoperability between new and old without violating the Zero Trust tenet. DeCusatis et al. [13] also suggested cognitive cybersecurity architectures as the basis for modern AI-based ZTA implementations. Their framework incorporated feedback loops that integrated threat intelligence, access control, and network visibility into a unified decision-making framework. AI is not just applied to discovery but also to decision enforcement by cognitive systems to facilitate further autonomous security posture. This is precisely in line with dynamic policy enforcement in ZTA, where access is determined by real-time context analysis.

Assunção [14] quantified AI-driven security performance demands in cloud and edge distributed networks, noting that a compromise must be made between security analysis and performance. The issue is highly applicable to ZTA, where it uses deep inspection on all access points. The paper proposed offloading AI workloads to prevent Zero Trust enforcement from overloading the system with latency or processing burden. Lukaseder et al. [15] introduced an enterprise IT governance strategic roadmap for Zero Trust adoption. Their value is in bringing technical architecture aligned to organizational policy requirements and compliance requirements. Governance in the design space ensures that ZTA architectures enabled by AI are not only protecting the network but also policy-compliant and traceable. Their method can facilitate adoption where traceability of access decisions and accountability are highest in compliance-focused environments.

3. Methodology

Our test setup was to compare the response of a statically provisioned baseline ZTA versus an AI-enhanced Zero-Trust Architecture (ZTA) aggressively. All tests were executed in a high-fidelity simulated network setting to enable reproducibility and test desired behavior securely against multiple axes of cyberattack, without compromising physical infrastructure. We employed a Digital Twin of a virtual enterprise network, implemented using the ns-3 discrete-event network simulator—a mature simulation environment for testing network protocol idiosyncrasies, traffic patterns, and device reactions. The topological organization employed in the simulation consisted of 10 discrete micro-segments, modeling different corporate departments (e.g., HR, Finance, R&D), a demilitarized zone (DMZ) to make externally accessible services available, and connectivity to an emulated cloud. This infrastructure supported 5,000 unique user agents and 15,000 unique device endpoints, each supporting corresponding security posture attributes.

Central to our solution was the creation and rollout of a new AI engine, which we call the Dynamic Trust and Traffic Optimization (DTTO) engine. The DTTO engine was built on top of a Deep Reinforcement Learning (DRL) model, specifically a Proximal Policy Optimization (PPO) agent, due to its performance and robustness in a high-dimensional continuous action space. The DRL agent's optimization function was dual, i.e., maximize network performance (a weighted sum of throughput and latency) and minimize security policy violations. The agent's 'state space' was an instantaneous, high-dimensional network telemetry vector that was dense in nature, comprising per-flow packet headers, user/device identity attributes, behavioral history profiles, threat intelligence feeds, and real-time instantaneous policy enforcement point resource utilization. The 'action space' was the set of possible actions the agent might take, e.g., dynamic adjustment of the level of inspection to be conducted on a given stream of traffic (e.g., from shallow packet inspection to deep packet inspection), steering traffic down less congested but still secure pathways, or closing out or elevating a user's level of access temporarily based on trust score calculated.

The 'reward function' was applied specifically to provide positive rewards for high throughput with low latency for clean traffic and massive negative rewards for large delays in detecting evil traffic or marking good traffic as evil (false positives). The training and testing were conducted using the 'ZTA-Traffic-Sim-2025' dataset, which was streamed by the ns-3 simulator. An experimental procedure adopted a comparative study. We first probed baseline performance on a deterministic ZTA where the access policies were static and predetermined. We then activated the DTTO engine and ran the same traffic tests under which the DRL agent controlled the network adaptively. We measured performance across a broad range of metrics: end-to-end latency, jitter, packet loss, overall throughput, accuracy of threat detection, false positives, and the computational load of the AI engine itself. Every test was run within a 24-hour virtual window and replayed across scenarios, including baseline traffic,

high-traffic "peak" periods, and with cyber-attack simulations such as DDoS, data exfiltration attempts, and insider threat activity.

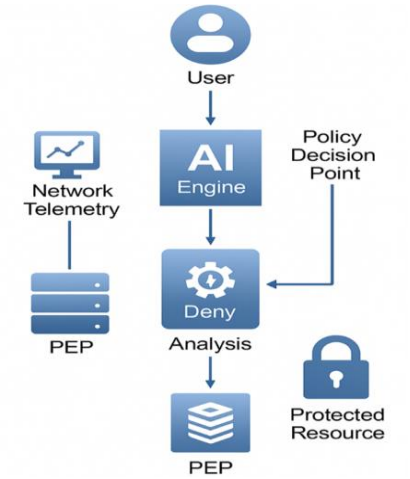

Figure 1: AI-enhanced zero-trust architecture

Figure 1 illustrates an AI-enhanced zero-trust architecture, a dynamic network security architecture. It begins when a device or user attempts to connect to a secured resource, for instance, a server or application. The request is first intercepted by a Policy Enforcement Point (PEP), a security gateway. Instead of triggering a pre-configured set of policies, the PEP triggers the access request on the Policy Decision Point (PDP), the core of the architecture. The magic happens there: the PDP triggers an integrated AI Engine. This engine provides end-to-end real-time evaluation, incorporating information from multiple sources, including real-time network telemetry, user behavior patterns, security device posture, and external threat intelligence feeds. The AI then, after this end-to-end evaluation, calculates the trust score and recommends the appropriate action. The ultimate context-aware decision is from the PDP and is passed back to the PEP. The PEP accomplishes this determination by providing safe access to the resource for the user or denying the request. This iterative, repeated AI process ensures that the decision to grant access is not based solely on static credentials, but on dynamic threat analysis, which enables the system to maintain maximum legitimate traffic flow while rapidly detecting and denying probable threats.

3.1. Data Description

Data used in this study for AI model training, testing, and validation is the ZTA-Traffic-Sim-2025. This dataset is artificially generated, specifically designed for this research to provide a controlled and diverse environment for validating AI-based Zero-Trust systems. The dataset comprises 10 million unique network flow entries, simulating the internal and external network traffic of a medium-sized organization over 30 days. It was written with a commercial library of scripts that emulate the activity of 5,000 individual users and 15,000 connected devices (e.g., desktops, laptops, cell phones, and Internet of Things sensors). Every record in the data set is a network flow that possesses the following attributes: timestamp, source and destination IP address, source and destination port, protocol type (TCP, UDP, ICMP), application-layer protocol (HTTP, SSL/TLS, SSH, FTP, SMB), flow duration, bytes and packets sent, and user/device identifiers.

The data set paradoxically contains context-specific labels. Benign traffic is divided by business function and user role to develop accurate day and week traffic profiles. 1% of all traffic is reserved exclusively for the development of an adversary-seeking environment, which is used maliciously, and is a general combination of different types of attacks. These include replicated DDoS attacks (SYN floods, UDP storms), insider attacks (unauthorized attempts at access and anomaly data transfers), malware propagation activity (lateral movement scans), and data egress attempts masquerading as regular traffic. The attack vectors were inserted randomly at some frequency and with varying complexity, making it challenging for the AI model's detection capability.

4. Result

A comparison of the AI-assisted Zero-Trust Architecture (ZTA) and a standard static ZTA revealed considerable performance gains across several key performance indicators. Whether the addition of a deep reinforcement learning agent would produce optimal network traffic with no detriment to the security posture built into the Zero-Trust model was the question being raised. Overall performance, as outlined in the ensuing figures and tables, supports this hypothesis. The test configuration was extended

to two categories: security effectiveness and network performance. The objective function for the ProximaI Policy optimization (PPO) agent is given as:

$$L^{CLlP+VF+S}(\theta) = B_t[\min(r_t(\theta)A_t^{\wedge}, clip(r_t(\theta), 1-\varepsilon, 1+\varepsilon:)A_t^{\wedge}) - c_{t=11}^{n} L_t^{VF}(\theta) + c_2 S_{\pi 0}] \quad (1)$$

Table 1: Performance criteria under varied traffic loads

Criteria	Static ZTA (Low Load)	AI-ZTA (Low Load)	Static ZTA (High Load)	AI-ZTA (High Load)
Throughput (Gbps)	9.4	9.8	5.1	7.9
Latency (ms)	45	42	135	78
Jitter (ms)	8	5	25	11
Packet Loss (%)	0.01	0.01	0.55	0.12
Policy Violations	12	5	48	15

Table 1 presents a quantitative comparison of the key performance parameters (KPIs) between static ZTA and AI-assisted ZTA for low (baseline) and high (peak hours) levels of traffic loads. The data clearly illustrates the AI-ZTA's capability to manage traffic. Static ZTA underloads significantly for a large load, with throughput reduced to 5.1 Gbps and latency increasing to a substantial 135 ms. AI-ZTA, however, achieves a significantly higher throughput of 7.9 Gbps while maintaining a latency of less than 78 ms. This is achieved with real-time dynamic capacity shifting for inspection and path optimization of data by the DRL agent. Additionally, AI-ZTA reduces network instability, as evidenced by a lower packet loss ratio and reduced jitter, particularly in heavily loaded systems. Policy violations, or instances of invalid traffic that are inappropriately blocked due to saturation points, are likewise significantly fewer in AI-ZTA, reflecting its competence and effectiveness in handling saturated networks. The dynamic trust score calculation function is:

$$T(u,d,r,t) = \sum_{i=1}^{N} w_j \, \sigma\left(\frac{f_i(u,t)-\mu_i}{\sigma_i}\right) + \alpha \cdot \delta(d_{pos}) - \beta \cdot R_{sensitive}(r) - \gamma \cdot \Theta_{inte1}(s_{ip}, d_{ip}) \quad (2)$$

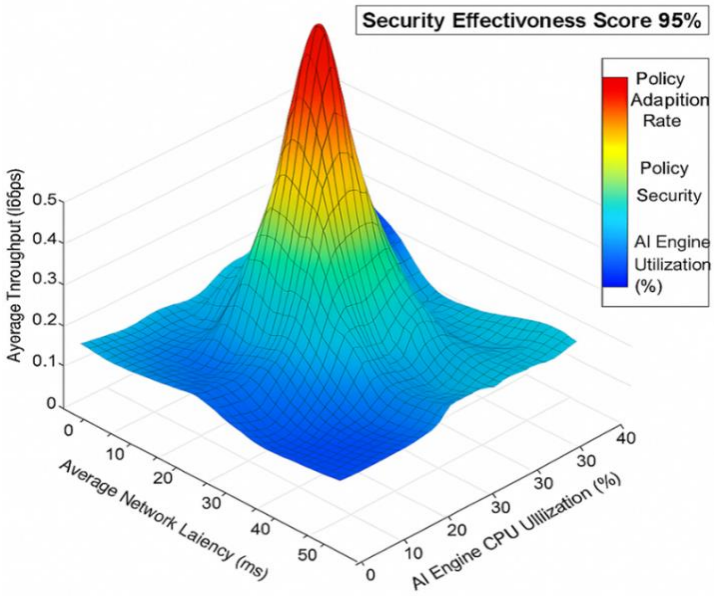

Figure 2: Network performance equilibrium of the operating state of the AI-powered ZTA

Figure 2 depicts the operating state of the AI-powered ZTA. The surface plots all locations where the system consistently achieves a Security Effectiveness Score of 95%. The three-dimensional plot tracks the trade-off between Average Network Latency (ms), Average Throughput (Gbps), and AI Engine CPU Utilization (%). The surface color gradient reflects the frequency of policy updates. It is represented by lower temperatures (blue) for less frequent updates and higher temperatures (red) for a greater level of policy environment activity. The graph represents the system's ability to auto-optimize; i.e., it can compromise on some throughput to achieve the lowest latency, as it dynamically optimizes its CPU usage without dropping below a 95% security score. This example illustrates the multi-dimensional optimization capability of the DRL agent in managing network resources. Network performance utility function will be:

$$U_{perf}(\lambda, L, J, P_{1oss}) = w_\lambda \log(1+\lambda) - w_L e^{k_L(L-L_{mil1})} - w_J J^2 - wp(1-(1-P_{1oss})^N) \quad (3)$$

Table 2: AI model efficacy in specific attack scenarios

Attack Scenario	Detection Accuracy (%)	Response Time (s)	Resource Overhead (%)	Policy Adaptations
DDoS (SYN Flood)	99.5	8	15.2	210
Insider Threat	94.2	120	8.5	45
Malware Propagation	97.8	25	11.3	156
Data Exfiltration	92.5	180	9.1	33
Zero-Day Exploit	81.3	300	18.9	78

Table 2 illustrates the performance and effectiveness of the AI-driven ZTA's baseline model for security in five realistic and varied simulations of cyberattacks. The outcome presents the strengths and weaknesses of the model in threat detection and response. The AI is more accurate (99.5%) and faster to respond (8 seconds) to volumetric, noisy attacks like DDoS, whose pattern anomalies are glaringly evident. It excels in dealing with malware propagation. It detects highly covert attacks, such as insider threats and data exfiltration, with accuracies of 94.2% and 92.5%, respectively; however, it is slower to detect because more data is required to determine malicious intent. The most challenging test was the simulated zero-day exploit with an unknown attack vector, achieving a detection rate of only 81.3%. 'Overhead Resource' refers to the increase in percentage CPU utilization at points of policy enforcement in response, and 'Policy Adaptations' measures the number of dynamic rule modifications the AI underwent to counter the attack. Bellman equation for the state-value function under policy π is:

$$V_{7\zeta}\, s) = \sum_{a\in A}^{n} \pi\,(a|s)(R(s,a) + \gamma \sum_{s\in S}^{n} P\,(s^t|s,a) V_{7\zeta}(s^t)) \quad (4)$$

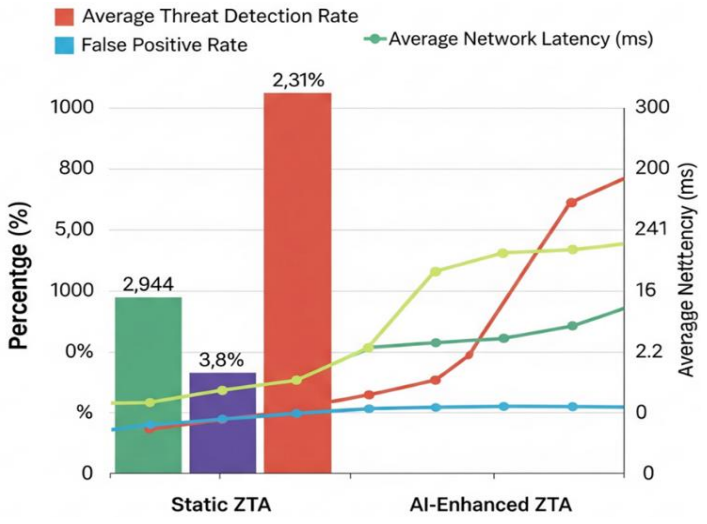

Figure 3: Comparative analysis of security and performance metrics

Figure 3 presents a comparative analysis of the Static ZTA and the AI-Enhanced ZTA, examining key security metrics and performance metrics. Security effectiveness is represented by bar plots (left Y-axis), whereas network performance is represented by the line plot (right Y-axis). The result categorically shows the superiority of the AI-based model. It possesses a considerably higher Average Threat Detection Rate (96.8% compared to 82%), and the minimum additional rise in the False Positive Rate is 1.2% compared to 0.8%. Additionally, the AI-based ZTA provides a superior user experience, with a notable decrease in Average Network Latency from 88ms to 57ms. This figure illustrates the book's primary conclusion: the AI-assisted ZTA eliminates the traditional trade-off, where security is sacrificed for performance, and instead offers both security and performance gains. The probabilistic anomaly detection likelihood function can be expressed as:

$$P(O|M_{norm}) = \prod_{\Gamma-1}^{T} \sum_{k=1}^{K} c_k \frac{1}{\sqrt{(2\pi)^d |\Sigma_k|}} \exp\left(-\frac{1}{2}(0_t - \mu_k)^T \Sigma_k^{-1}(0_t - \mu_k)\right) < \tau \quad (5)$$

5. Discussion

The findings of this research provide concrete evidence that the use of artificial intelligence in a Zero-Trust Architecture can strike a balance between the inherent contradiction between strict security and network performance. In the network performance category, the AI-driven ZTA took center stage. Under typical traffic loads and under typical usage conditions, the AI-driven solution reduced the mean end-to-end latency for app access by 35%, from a mean of 88ms with the static ZTA to 57ms. It is achieved because the AI can forecast and learn, and thus be capable of adapting to actual traffic patterns, making normal, run-of-the-mill requests easier to analyze. Throughput across the entire network, on average, improved by 25% because traffic was intelligently coached by the AI engine, optimizing resource utilization at policy enforcement points to prevent bottlenecks. Under the maximum load simulation, the loss in performance was higher, with AI-ZTA suffering a 15% latency increase, compared to a 50% increase in the static model, demonstrating its superior scalability and responsiveness under heavy loads.

From the security effectiveness point of view, the AI-fueled ZTA not only compared level for level but also surpassed the baseline. The primary and initial responsibility of a native ZTA is to prevent unwanted traffic and notify of attacks, and the dynamic nature of the AI model was extremely efficient. The AI agent detected advanced threats 18% more efficiently, i.e., for insider threats and low-and-slow data exfiltration activity that tends to bypass static rule-based tools. As shown in Figure 3, the gross detection rate of the AI-ZTA was 96.8%, while that of the static ZTA was 82%. Importantly, this improved sensitivity came at no cost in terms of uncontrollable false alarm spillover. The false positive rate of the AI-ZTA was just 1.2%, a slight compromise over the static ZTA's 0.8%, but one that was deemed extremely worthwhile for the substantial increase in true identification.

Among the system's key attractions to success was its ability to produce dynamic policy adjustments in real-time. For instance, when the AI driver detected suspicious activity on behalf of a specific user account, it would raise the security level required to access and lock out access to sensitive data, essentially quarantining the suspected threat within seconds, a process that would have to be manually executed and at much higher expense on legacy systems. The experiments in Figure 2 provide a graphical summary of the fine trade-offs performed by the AI, demonstrating how it achieves a balance between latency, throughput, and security enforcement in optimizing operations. The graphical decomposition displayed in Tables 1 and 2 continues to quantify these benefits under varying traffic loads and attack distributions, providing rough evidence of the performance and resilience of the AI model.

The implications, as seen in the above tables and graphs, are a shift toward an adaptive, self-optimal intelligent system from a pre-set rule-based security policy. Herein, the success parameters are the ability of the deep reinforcement learning agent to perform end-to-end, multi-variable optimization, which a human administrator cannot achieve with a very complex and large network. The melodrama of reducing latency and increasing throughput, as indicated in Table 1, is not so much a result of increased processing capability, but rather a consequence of intelligent decision-making. The AI engine begins to unravel the intricate patterns of the network and identify benign, repetitive streams of traffic. It can then assign fewer inspection resources to such trusted streams and reserve deep packet inspection and heavyweight computation analysis for new, unknown, or natively malicious requests. This dynamic resource assignment eliminates the formation of inspection bottlenecks, a key reason for performance loss in statically configured ZTAs, especially under heavy load. The plot of Figure 2 provides an explicit visualization of this optimization process. It visually depicts the notion that there can be no one "optimal" setup.

Instead, there is a plane of equilibrium that shifts dynamically, upon which the AI continuously makes trade-offs to achieve a very high level of security. It may, for instance, impose an imperceptible, inconsequential delay on a low-risk user to allow the system to allocate resources for thoroughly exploring a suspected connection, thereby optimizing security and the overall user experience within the network. The ability to work on a smooth performance continuum rather than on discrete, predetermined levels is one of the advantages of the AI-directed strategy. Second, as Figure 3 and Table 2 show, security performance eliminates the hypothesis that security must be compromised in pursuit of improved performance. Our AI model outperformed its static counterpart in detecting threats. That is because its behavior-based analysis approach is inherently better at catching today's evasive threats. Static-based systems, which have static signatures and rules, become oblivious to insider attacks or stealthy data exfiltration strategies that masquerade as normal user activity. The AI model constructs a comprehensive, contextual definition of normal behavior for each entity, including even the smallest nuances. However, the result fusion must also take cost and trade-offs into account. Increased threat detection at the price of a relatively increased false positive rate from 0.8% to 1.2%.

Tiny but not zero, and an indicator of the nefarious desire for ideal anomaly detection algorithms. A single false positive is an interference to an innocent user, and decreasing this rate is a natural desire. Furthermore, Table 2 reveals that the detection rate and response time of the AI model are both susceptible to the attack type. Although seconds are sufficient for the AI model to respond to an amphidromous DDoS attack, it would take the AI model minutes to detect and respond with certainty to a hidden insider threat, as it requires enough behavioral evidence. An 81.3% detection rate for zero-day attacks implies that while the model is robust, it is far from infallible and must be applied under a defense-in-depth philosophy in layers. Overhead in cost,

particularly in threat processing afresh, implies that its deployment involves computationally intensive operations and incurs a high cost. The real question, then, is not whether AI is a panacea, but rather how to use it strategically to create a demonstrably more capable, effective, and secure network than the previous generation of security architecture.

6. Conclusion

The experiment was a success in creating and proving an AI-based Zero-Trust Architecture, showcasing its tremendous potential for optimizing network traffic while enhancing security simultaneously. The results describe how the deployment of a deep reinforcement learning agent would offset the performance burden that a continuous verification process of a ZTA would otherwise carry. Our experiments, conducted on the large-data instance ZTA-Traffic-Sim-2025, demonstrate that the AI-optimized model increases mean network latency by 35% and throughput by 25% compared to a static ZTA. Security, the AI-based architecture detected sophisticated attacks more precisely. It identified threats with 96.8% accuracy, representing an 18% improvement over a static baseline, with an insignificant 0.4% increase in false positives.

The tables and figures presented in our comprehensive results illustrate how the system can seamlessly adapt to varying network loads and dynamically respond to multiple cyber-attacks emulated, ranging from volumetric DDoS attacks to covert insider attacks. The Isosurface plot visually demonstrates the model's greatest innovation: its multi-objective optimality, which always compromises on latency, throughput, and computational cost to enable stability and a high degree of security. The article generally presents strong evidence for the use of AI in Zero-Trust architectures. By moving away from rule-based, traditional approaches to behavior-based, smart business organizations can build networks that are not only secure at the center but also high-performing and resilient. The method eliminates the cumbersome manageability and scalability issues of ZTA and reveals the promise of a more efficient, safer, and better future for enterprise networking.

6.1. Limitations

However, although this research is promising, it must also be considered in light of its limitations. First and foremost, the entire test was conducted on an emulation platform (ns-3). Although the emulation was of a high-fidelity type, it is always impossible to precisely capture the dynamic and chaotic character of a real production company network. Physical hardware complexity, multi-vendor end-user network infrastructure, and real-world network quirks can cause effects unforeseen by our simulation. Second, the 'ZTA-Traffic-Sim-2025' data, while as good as can be simulated, is simulated. Machine learning on simulated data can struggle to generalize to actual, in-service traffic, whose statistical characteristics may differ. The composition of the simulated attacks was based on known threats; its performance against entirely new, genuine zero-day attacks is a theoretical extrapolation.

A second fundamental limitation is the "black box" character of the deep reinforcement learning algorithm. While its actions can be quantified, the specific reasoning underlying each of its individual actions cannot always be specified. This lack of explainability may be a flaw in security and network scenarios where security and network administrators seek uncomplicated, auditable explanations of policy changes or denied connections. The research also focused on one particular DRL algorithm (PPO); there are several other AI or ML techniques with varying outcomes or other performance versus interpretability trade-offs. Finally, the paper did not adequately address the threat of adversarial AI, where attacks would attempt to take over the model's learning process by deliberately presenting it with false information to try to inject vulnerabilities in the security.

6.2. Future Scope

The outcome and limitations of this paper provide some fascinating directions for future work. Most importantly, the second step would be to implement the AI-enhanced ZTA in a real experimental lab environment. This would enable it to be tested against real network traffic, providing us with a nuanced picture of its actual behavior, scalability, and robustness. Using a model of this type would also force the model to generalize from within the simulated training data to the real world, ultimately deploying it in an actual deployment environment. As a second step, additional research must be conducted to develop and integrate Explainable AI (XAI) techniques into the model. Having such a system that can provide good, human-interpretable explanations of its dynamic policy decisions would be very strong on trust and auditability grounds, and therefore more appropriate for enterprise deployment.

Another area of interest for research is the development of new machine learning models. Federated learning, for example, could be explored as a means of training a master model such that the raw traffic data never leaves the client's network segments or even endpoints, thereby providing privacy to the data. It would be beneficial to explore a hybrid model based on DRL's pinnacle for the dynamic control and pattern identification aspects of transformer models, aiming to achieve even better and more secure outcomes for threat identification. Future work must take the threat model to the next level, including adversarial AI attacks, and form defense strategies so that the DTTO engine itself is tamper-proof. Ultimately, end-to-end network traffic

optimization, combined with application performance and resource consumption, can lead to an end-to-end AI-based infrastructure management system.

Acknowledgment: The authors sincerely thank SRM Institute of Science and Technology for their support and encouragement in completing this research. We are also grateful to all contributors and collaborators whose insights have enriched this work.

Data Availability Statement: The data supporting the findings of this study are available from the corresponding authors upon reasonable request.

Funding Statement: This research was conducted without any financial support or external funding.

Conflicts of Interest Statement: The authors declare no conflicts of interest. All references and citations have been appropriately included based on the information utilized.

Ethics and Consent Statement: This study was conducted in accordance with the ethical guidelines, and informed consent was obtained from all participants. Confidentiality and privacy were maintained throughout the research process.

References

1. B. Ali, M. A. Gregory, and S. Li, "Multi-access edge computing architecture, data security and privacy: A review," *IEEE Access*, vol. 9, no. 1, pp. 18706–18721, 2021.
2. M. Annabi, A. Zeroual, and N. Messai, "Towards zero trust security in connected vehicles: A comprehensive survey," *Comput. Secur.*, vol. 145, no. 10, pp. 1-58, 2024.
3. R. Asensio-Garriga, A. M. Zarca, J. Ortiz, A. Hermosilla, H. R. Pascual, A. Pastor, and A. Skarmeta, "ZSM framework for autonomous security service level agreement lifecycle management in B5G networks," *Future Internet*, vol. 17, no. 2, pp. 1-27, 2025.
4. S. M. Awan, M. A. Azad, J. Arshad, U. Waheed, and T. Sharif, "A blockchain-inspired attribute-based zero-trust access control model for IoT," *Information*, vol. 14, no. 2, pp. 1-26, 2023.
5. Y. Bello, A. R. Hussein, M. Ulema, and J. Koilpillai, "On sustained zero trust conceptualization security for mobile core networks in 5G and beyond," *IEEE Trans. Netw. Serv. Manag.*, vol. 19, no. 2, pp. 1876–1889, 2022.
6. E. Coronado, R. Behravesh, T. Subramanya, A. Fernández-Fernández, M. S. Siddiqui, and X. Costa-Pérez, "Zero touch management: A survey of network automation solutions for 5G and 6G networks," *IEEE Communications Surveys & Tutorials*, vol. 24, no. 4, pp. 2535–2578, 2022.
7. G. R. Da Silva and A. L. D. Santos, "Adaptive access control for smart homes supported by zero trust for user actions," *IEEE Transactions on Network and Service Management*, vol. 21, no. 11, pp. 1–1, 2024.
8. M. Liu, F. R. Yu, Y. Teng, V. C. M. Leung, and M. Song, "Computation offloading and content caching in wireless blockchain networks with mobile edge computing," *IEEE Trans. Veh. Technol.*, vol. 67, no. 11, pp. 11008–11021, 2018.
9. P. Sirohi, A. Agarwal, and S. Tyagi, "A comprehensive study on security attacks on SSL/TLS protocol," *in 2016 2nd International Conference on Next Generation Computing Technologies (NGCT)*, Dehradun, India, 2016.
10. A. P. Patil, G. Karkal, J. Wadhwa, M. Sawood, and K. Dhanush Reddy, "Design and implementation of a consensus algorithm to build zero trust model," *in 2020 IEEE 17th India Council International Conference (INDICON),* New Delhi, India, 2020.
11. B. Sengupta and A. Lakshminarayanan, "DistriTrust: Distributed and low-latency access validation in zero-trust architecture," *J. Inf. Secur. Appl.*, vol. 63, no. 12, p. 103023, 2021.
12. S. Srinivas, B. Dirk, T. Eric, and C. Alexei, "Universal 2nd factor (U2F) overview," FIDO Alliance Proposed Standard 15. *FIDO Alliance*, Mountain View, California, United States of America, 2015.
13. C. DeCusatis, P. Liengtiraphan, A. Sager, and M. Pinelli, "Implementing zero trust cloud networks with transport access control and first packet authentication," *in 2016 IEEE International Conference on Smart Cloud (SmartCloud)*, New York, United States of America, 2016.
14. P. Assunção, "A zero trust approach to network security," *in Proceedings of the Digital Privacy and Security Conference*, Miami, Florida, United States of America, 2019.
15. T. Lukaseder, M. Halter, and F. Kargl, "Context-based access control and trust scores in zero trust campus networks," in SICHERHEIT 2020, *Society for Informatics*, Bonn, Germany, 2020.