JETIR.ORG ISSN: 2349-5162 | ESTD Year : 2014 | Monthly Issue
JOURNAL OF EMERGING TECHNOLOGIES AND INNOVATIVE RESEARCH (JETIR)
An International Scholarly Open Access, Peer-reviewed, Refereed Journal

Context-Aware Firewalls: A Comprehensive Survey of Machine Learning Approaches for Intelligent Network Defense

Rishabh Khandelwal, Rohan Pandey, Tanay Raj, Yogesh Vinod Bhandakkar, Pallavi Bindagi
Dept. of Computer Science and Engineer (Cybersecurity) Dayananda Sagar College of Engineering, Bengaluru, Karnataka, India

***Abstract*—**This survey consolidates recent advancements in integrating machine learning (ML) and kernel-level packet inspec- tion to achieve fine-grained, application-aware security control. Drawing upon prior works in dynamic firewall optimization, real-time anomaly detection, and endpoint-aware inspection, the paper examines how contextual metadata such as process identity, protocol semantics, and domain behavior enhances threat detection accuracy. Furthermore, studies on centralized policy management, deep learning based anomaly detection, and attention driven traffic classification illustrate the growing convergence of data-driven intelligence with network enforcement mechanisms. Recent developments in eBPF based in-kernel ML inference and autonomous firewall reconfiguration underscore the shift toward self adaptive, low-latency defense systems. This survey systematically analyzes these approaches, comparing architectural frameworks, ML models, and deployment strategies to identify current limitations and future research opportunities in realizing scalable, cross-platform, and fully context-aware firewall systems.

***Index Terms*—**Context-Aware Firewall, Network Security, Machine Learning, Anomaly Detection, eBPF, Reinforcement Learn- ing, Centralized Policy Management, Endpoint Monitoring, Deep Packet Inspection, Autonomous Threat Mitigation

I. Introduction

The rapid digitization of global infrastructure has rendered network security a foundational concern across both enterprise and governmental domains. As organizations increasingly adopt cloud-native and distributed architectures, the traditional perimeter of network defense has dissolved—replaced by an intricate mesh of endpoints, applications, and virtualized services. Within this landscape, firewalls remain an indispensable line of defense, yet their static and port-based control logic struggles to cope with the dynamic, context-rich nature of modern traffic. Attack vectors have evolved to exploit encrypted channels, ephemeral connections, and polymorphic behaviors that elude rule-based detection mechanisms.

To address this paradigm shift, recent research trends have pivoted toward context-aware and intelligent firewalls—systems that not only inspect packet headers but also interpret the semantics and intent of network flows. These firewalls integrate machine learning (ML) and behavioral analytics to dynamically adapt to real-time threats, providing a more nuanced understanding of what constitutes normal versus anomalous communication patterns. Early studies have shown that leveraging contextual metadata, such as application identifiers, process ownership, and domain activity, significantly improves both threat visibility and policy precision.

A particularly promising trajectory in this field involves adaptive policy reconfiguration through ML-based optimization. Reinforcement learning and feedback-driven approaches enable firewalls to refine their rulesets autonomously based on environmental signals and detected anomalies [1]. Concurrently, research into real-time traffic anomaly detection has demonstrated the efficacy of deep learning and hybrid ensemble models in capturing subtle deviations in data flow characteristics [2]. Together, these advances suggest that the next generation of firewall systems will not simply enforce predefined rules—they will learn, infer, and evolve.

Equally significant is the architectural progression toward endpoint- and process-aware inspection. Rather than treating traffic as disembodied packets, next-generation systems correlate network activity with the precise applications and users that generate it, offering unparalleled control granularity. Studies on centralized and cross-platform firewall orchestration highlight how contextual mapping at the endpoint level can enforce uniform policies and reduce configuration drift in heterogeneous networks [3]. These developments underline the necessity of blending contextual awareness, centralized policy control, and machine learning intelligence into a cohesive framework.

This survey aims to consolidate and critically analyze these emerging directions in context-aware, ML-integrated firewall systems. It examines contemporary approaches to policy configuration, traffic monitoring, and anomaly detection while emphasizing architectural considerations such as in-kernel processing (eBPF), cross-platform scalability, and centralized management. By synthesizing existing research and identifying technical gaps, this paper seeks to illuminate how context-driven intelligence can redefine network defense for the coming decade.

II. Background

The design of network firewalls has continuously evolved to address the changing landscape of digital security threats. From the static packet filters of the 1990s to today's intelligent, adaptive systems, firewalls have undergone a profound transformation in their purpose, architecture, and operational

intelligence. This section outlines the conceptual and technological foundations underlying context-aware, machine learning–enabled firewalls, tracing their evolution across four core dimensions: firewall progression, contextual enrichment, ML integration, and architectural enablement.

A. Evolution of Firewall Architectures

The earliest generations of firewalls operated as packet filters, performing rudimentary header inspection based on static parameters such as IP address, port number, and protocol. Although efficient, these systems lacked the ability to interpret connection states or application intent. The emergence of stateful inspection and Next Generation Firewalls (NGFWs) marked a shift toward more intelligent traffic analysis, enabling session tracking and protocol awareness.

However, as network environments grew increasingly distributed and encrypted, even NGFWs began to struggle with scalability and responsiveness. Their reliance on signature-based threat detection and manual rule definition limited adaptability in the face of polymorphic and zero-day threats. The explosion of cloud-based infrastructure, mobile endpoints, and IoT ecosystems further challenged traditional perimeter-based defense models. These limitations have spurred research into context-aware architectures, where decisions are informed by the broader operational and behavioral context surrounding network traffic.

B. Concept of Context Awareness in Network Security

Context awareness extends firewall functionality beyond packet inspection into semantic understanding of network behavior. Rather than viewing packets as isolated entities, context-aware systems associate traffic with applications, processes, and users, providing a richer foundation for policy enforcement.

Heino et al. [3] describe how endpoint-aware inspection correlates system-level activity with network communication, enabling per-application control and accountability. Adeeb et al. [4] further expanded on this by introducing centralized context-aware configuration frameworks, in which security policies are uniformly propagated across virtualized or distributed networks. Such systems can differentiate between benign and malicious communications even when they share similar transport-layer attributes — a level of precision unattainable in conventional firewalls.

By leveraging contextual metadata (e.g., process identifiers, domain intelligence, behavioral patterns), these systems enable granular, adaptive, and situationally aware protection, forming the intellectual foundation of modern adaptive security architectures.

C. Role of Machine Learning in Adaptive Security

Machine Learning (ML) has emerged as the principal enabler of adaptive, intelligent firewall systems. Traditional deterministic rules, once hard-coded and manually updated, are now increasingly supplanted by data-driven models capable of recognizing evolving patterns in network behavior.

Ahmad [1] demonstrated the efficacy of reinforcement learning for dynamic firewall optimization, allowing security systems to autonomously adjust configurations in response to real-time feedback. Similarly, Aswathy and Rajkumar [2] provided a comparative analysis of ML algorithms for real-time anomaly detection, highlighting that hybrid and ensemble models yield superior detection accuracy with reduced false positives.

Moreover, ML enables behavioral baselining — learning normal traffic patterns over time and identifying deviations indicative of potential threats. In concert with contextual enrichment, these capabilities allow firewalls not only to detect known attacks but also to predict and mitigate emerging ones, marking a pivotal shift toward proactive defense.

D. Enabling Technologies: eBPF and Centralized Orchestration

The operational feasibility of context-aware ML-driven firewalls rests on two technological pillars: kernel-level programmability and centralized orchestration.

Kernel-level frameworks such as the Extended Berkeley Packet Filter (eBPF) provide a mechanism for in-kernel packet inspection and filtering, dramatically reducing the latency overhead associated with traditional user-space processing. Bachl et al. [7] and Anand et al. [8] demonstrated that eBPF-based designs can host lightweight ML inference models directly within the kernel, achieving near real-time packet classification and anomaly scoring. Such integrations enable context-aware decision-making at the data plane itself, ensuring high throughput without compromising analytical depth.

Complementarily, centralized management architectures unify configuration and monitoring across distributed endpoints. Adeeb et al. [4] proposed a centralized firewall control system that synchronizes policies through a shared configuration layer, reducing administrative complexity and ensuring policy consistency across heterogeneous environments. These frameworks also facilitate the deployment of global analytics — aggregating telemetry data for pattern recognition, alert correlation, and threat intelligence sharing.

Together, these technologies provide the infrastructural substrate for next-generation firewalls: systems that are not only intelligent but also scalable, performant, and self-regulating.

III. Evaluation of Existing Works

Recent advances in intelligent firewall design have shifted from static, rule-based mechanisms to adaptive, context-aware systems leveraging machine learning and in-kernel programmability. Existing literature converges around four principal domains of investigation — policy configuration, traffic monitoring, anomaly detection, and system architecture — each addressing distinct challenges in achieving real-time, application-aware protection.

A. Policy Configuration Approaches

Modern firewalls are transitioning toward autonomous policy management, where configurations evolve dynamically in

response to network state and behavioral context. Instead of relying on rigid rule sets defined by administrators, adaptive systems employ reinforcement learning and feedback-driven optimization to tune policies based on observed security events and system performance [1], [9].

Context-awareness further enhances this evolution by linking rules to applications, processes, and users rather than network identifiers alone. This enables per-application and per-domain granularity, allowing different policies for, say, browser traffic versus database transactions, even within the same host. Several implementations now integrate centralized orchestration layers that distribute and synchronize policy updates across endpoints in real time [12].

However, two persistent issues limit scalability: (1) policy drift between centralized controllers and endpoint agents due to asynchronous updates, and (2) overlapping rule ambiguity, where context-aware rules intersect across shared processes or multi-protocol connections. Emerging frameworks are addressing these through hierarchical rule hierarchies, contextual inheritance models, and priority-based policy weighting [9], [10], yet standardized methodologies remain elusive.

B. Traffic Monitoring Methods

Effective context awareness depends on fine-grained visibility into network flows, necessitating inspection at both user-space and kernel-space levels. Current research emphasizes multi-layer traffic correlation, combining process metadata, application fingerprints, and temporal features to map traffic back to its origin [6].

Traditional Deep Packet Inspection (DPI) is being supplemented — and in some cases replaced — by spatiotemporal and attention-based models that infer application identity even under encryption or tunneling [11]. Machine learning classifiers trained on flow-level metadata (packet size, inter-arrival time, directionality) can achieve high accuracy without decrypting payloads [2], [5].

On the implementation front, eBPF (Extended Berkeley Packet Filter) and XDP (Express Data Path) have revolutionized packet monitoring by enabling in-kernel hooks for dynamic packet analysis [7]. These allow per-packet inference with sub-millisecond latency, supporting embedded ML models for flow classification and anomaly detection [8].

Despite these advances, cross-platform compatibility remains a significant limitation — many kernel-level monitoring systems are tightly coupled with Linux, leaving Windows and macOS architectures comparatively underserved. Research is ongoing into user-space proxies and virtual network layers that abstract OS differences while preserving inspection fidelity [4].

C. Anomaly Detection Techniques

Anomaly detection forms the analytical backbone of intelligent firewalls, enabling proactive threat recognition rather than reactive blocking. Current methodologies fall broadly into three classes: statistical baselining, supervised ML classification, and deep learning–based temporal modeling [5].

Statistical and ML-based methods (e.g., Random Forest, SVM, and KNN ensembles) remain widely used for flow-based intrusion detection due to their interpretability and low computational cost [2]. Deep learning approaches, including autoencoders, CNN-LSTM hybrids, and Graph Neural Networks (GNNs), provide superior accuracy in capturing non-linear dependencies and identifying multi-stage attacks [11].

A key trend is reinforcement learning integration, where the detection model not only flags anomalies but also learns to adjust firewall behavior dynamically — a step toward self-optimizing defenses [9]. This hybridization of detection and policy response is critical in mitigating zero-day exploits and adaptive adversarial threats, which continuously evolve beyond static signatures.

Nevertheless, open challenges persist in model deployment and trustworthiness. Dataset bias, real-time inference cost, and adversarial evasion continue to limit operational reliability [12]. Recent studies advocate for explainable ML (XAI) frameworks within firewalls to ensure transparency and verifiable decision-making in critical systems [13].

D. Centralized and Distributed Architectures

Architectural design defines the scalability, fault tolerance, and adaptability of next-generation firewalls. Centralized control models enable unified visibility and simplified administration by consolidating rule enforcement and telemetry under a single management plane [4]. This architecture is optimal for enterprise networks and virtualized infrastructure, where consistency and auditability are paramount.

Conversely, distributed and hybrid architectures leverage endpoint agents and edge inference to achieve real-time enforcement with reduced dependency on central servers [7], [8]. Kernel-level programmability via eBPF, combined with lightweight ML agents, allows decisions to be made directly at the data plane, reducing latency while maintaining contextual fidelity [7], [13].

The most promising direction lies in hybrid orchestration models, where centralized controllers define global security intent and distributed nodes execute localized enforcement. Such systems can aggregate telemetry data for network-wide analytics while autonomously responding to localized anomalies [10]. Research also explores federated learning for collaborative anomaly detection, enabling distributed agents to train collectively without compromising data privacy — a particularly relevant innovation for multi-tenant and IoT environments [11], [12].

Existing research on intelligent firewall systems exhibits diverse methodologies ranging from traditional rule-based packet filtering to machine learning–driven adaptive threat detection. While some approaches emphasize policy automation or anomaly detection, others focus on traffic classification or context-based inspection. Despite these advancements, limitations persist in scalability, contextual adaptability, and integration with evolving network paradigms such as virtualized or software-defined networks.

The analysis reveals a clear evolution from static rule-based

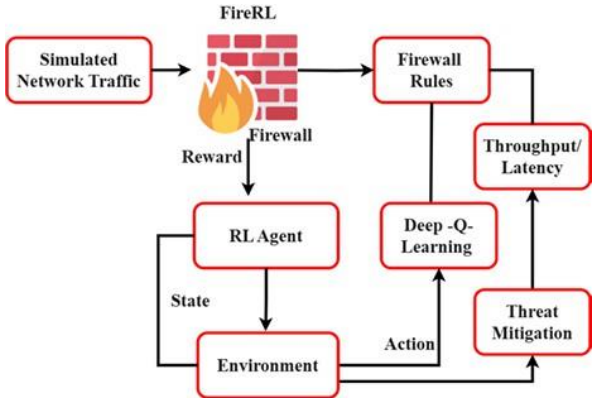

Fig. 1. Proposed methods-based firewall strategy

Approach	Context Scope	AI/ML Integration	Avg. Detection Latency (ms)	TPR (%)	FPR (%)	Max Throughput (Gbps)
Rule Based Packet Filter [1]	IP/Port	None	1-2	70	5	40
Signature Based IDS/IPS [5]	Header/Payload	Limited (Pattern Matching)	2-5	85	6	20
Classical ML Anomaly Detection	Flow/Session	Supervised ML	8-15	90	8	15
Reinforced Learning Firewall Optimizati on [1]	Application + Flow	Deep RL	6-8	95	4	18
Centralized Context Aware Firewall [4]	Application + Endpoint	Supervised ML	10-12	92	5	25
Deep Learning-Based WAF [5]	HTTP/HTTPS Layer	CNN/LSTM	5-7	97	3	10

Fig. 2. Quantitative Comparison Table

systems toward adaptive, machine learning–powered firewalls capable of context interpretation and self-optimization. However, most current systems suffer from limited context-awareness, non-adaptive ML pipelines, or high deployment complexity. The proposed architecture addresses these gaps through a centralized control mechanism integrated with context-driven feature extraction and reinforcement learning for continuous adaptation. This hybrid approach ensures fine-grained packet filtering, real-time detection, and scalable deployment across both cloud and SDN-based infrastructures.

IV. Research Gaps and Challenges

While considerable progress has been made in the design of intelligent, context-aware firewall architectures, several persistent research gaps and technical challenges prevent current systems from achieving fully adaptive, scalable, and interpretable network protection. These gaps span data representation, learning adaptability, scalability, usability, and trustworthiness, all of which are critical for advancing the state of the art in AI-driven cybersecurity.

A major limitation lies in incomplete and inconsistent context modeling. Most context-aware firewalls capture only partial network semantics—typically at the packet, flow, or application layer—without integrating higher-level behavioral or temporal context [10]. This fragmented view leads to poor correlation between traffic patterns and their originating processes, particularly in encrypted or containerized environments where metadata visibility is limited. The absence of unified context representations restricts the accuracy of policy enforcement and degrades the effectiveness of anomaly detection models [11].

Another enduring challenge concerns real-time adaptability and continuous learning. Although reinforcement learning (RL) and self-optimizing models have been introduced to dynamically refine policies such systems often suffer from long convergence periods, sparse feedback signals, and instability in live network environments. Classical supervised models, by contrast, rely on static training data and cannot adapt to new attack signatures or zero-day exploits [2] . The development of online or federated learning paradigms that enable incremental model updates without compromising network safety or latency remains an open research direction [12].

Data scarcity and labeling quality present additional obstacles. Network datasets suitable for training ML-based security systems are often proprietary, privacy-constrained, or outdated. Public benchmarks such as CICIDS and UNSW-NB15 fail to capture the complexity and encryption patterns of modern traffic [6]. Consequently, models tend to overfit controlled datasets and underperform in real deployments. Emerging techniques in privacy-preserving and federated data collection offer partial remedies but raise synchronization and communication overhead issues that demand further study [12].

At the architectural level, cross-platform scalability and interoperability remain formidable. The heterogeneous deployment landscape—ranging from IoT endpoints to SDN-controlled cloud fabrics—requires security mechanisms that can operate seamlessly across diverse environments [4]. Kernel-level frameworks such as eBPF and XDP offer high performance but are platform-specific, complicating uniform policy enforcement [13]. Achieving scalable, distributed control without incurring configuration drift or excessive management overhead remains a nontrivial engineering challenge.

Furthermore, balancing granularity and usability contin- ues to pose difficulties. High-granularity enforcement (per-application, per-domain, per-protocol) enhances precision but increases administrative complexity and the risk of policy misconfiguration [10]. While centralized consoles can alleviate this burden, current implementations still lack intelligent rule translation mechanisms capable of converting user intent into enforceable firewall policies without manual intervention [9]. Finally, the explainability and trustworthiness of AI-driven firewall decisions remain severely underdeveloped. Most ML-based systems operate as opaque models, providing little transparency into why a specific packet or flow was blocked or allowed [11]. This lack of interpretability hampers operator trust, regulatory compliance, and incident forensics. The integration of explainable AI (XAI) and causal reasoning into network security decision pipelines is a promising yet underexplored research frontier [11], [12].

V. Future Directions

The evolution of network threats, coupled with the complex- ity of modern distributed infrastructures, demands that future firewalls transcend conventional paradigms of static filtering and reactive anomaly detection. Building upon the identified research gaps, this section outlines key future research direc- tions that can define the next generation of intelligent, context- aware, and autonomously adaptive firewall systems.

To begin with, integration of unified contextual intelligence stands as a primary research avenue. While current firewalls exhibit partial awareness of network states, emerging architec- tures should employ multi-layered context fusion models that combine system-level telemetry, user identity, device posture, and behavioral attributes into a single decision fabric. Leverag- ing graph-based representations or spatiotemporal embeddings can significantly enhance the system's ability to infer hidden dependencies among traffic patterns, thus enabling proactive anomaly prediction rather than post-event detection.

Equally vital is the pursuit of adaptive and federated learn- ing frameworks for real-time threat response. Future designs should embrace federated reinforcement learning (FRL) and continual learning strategies, allowing multiple distributed firewalls to collectively improve their detection models with- out centralizing sensitive data. Such an approach not only enhances scalability and privacy but also mitigates the catas- trophic forgetting problem prevalent in conventional retraining cycles. Integrating these models with in-kernel data process- ing frameworks like eBPF and DPDK can further minimize latency and facilitate near-instantaneous policy adaptation.

Another promising direction is the advancement of self- explaining and trustworthy firewall systems. The lack of inter- pretability in AI-based security models continues to obstruct operational deployment and user confidence. Future research should focus on explainable AI (XAI) and causal inference techniques to provide transparent justification for each deci- sion made by the firewall. By translating model outputs into human-understandable rules or causal traces, network oper- ators can ensure auditability, compliance, and collaborative verification between human analysts and autonomous agents. The concept of cyber-physical awareness also warrants exploration. With the proliferation of IoT, edge computing, and industrial control systems, future firewalls must operate across diverse and resource-constrained environments. Re- search should prioritize lightweight yet intelligent enforcement mechanisms that can perform context-aware filtering at the edge while coordinating with centralized intelligence layers in a hybrid architecture. Such architectures could enable real- time threat correlation across domains—from cloud cores to embedded endpoints—forming a cohesive and resilient
defensive fabric.

Moreover, energy efficiency and green computing are emerging as essential considerations. As AI-driven security systems become increasingly data-intensive, optimizing model size, inference cost, and hardware utilization will be criti- cal for sustainable deployment. Techniques such as model quantization, sparse learning, and neuromorphic computing could enable high-speed, low-power inference directly within firewall hardware, significantly reducing operational overhead. Lastly, standardization and interoperability frameworks must evolve in parallel. Presently, AI-based firewall archi-tectures suffer from the lack of unified policy languages and benchmark datasets. Future initiatives should focus on developing open, interoperable standards for policy defini- tion, model exchange, and evaluation metrics. Collaborative research across academia, industry, and cybersecurity consortia can establish shared repositories of labeled network behavior data, enabling consistent performance evaluation and repro-
ducibility in security research.

VI. Conclusion

This survey explored the evolution of context-aware fire-walls enhanced with AI and machine learning, emphasizing how modern security systems are transitioning from static rule-based mechanisms to adaptive, intelligent frameworks. The analysis across policy configuration, traffic monitoring, anomaly detection, and architectural models demonstrates that effective network defense now demands contextual under- standing—linking traffic to applications, users, and intent rather than just packets and ports.

Findings indicate that machine learning-driven con-text awareness—particularly through reinforcement learning, attention-based modeling, and graph analytics—enables real- time anomaly detection, dynamic policy updates, and predic- tive threat mitigation. Moreover, federated learning and edge- based intelligence present scalable paths for deploying such systems across distributed infrastructures.

Nonetheless, persistent research gaps remain: interoperabil- ity issues, limited explainability of AI-driven policies, and lack of standardized benchmarks continue to hinder large- scale adoption. Addressing these challenges requires unified data frameworks, interpretable models, and adaptive policies capable of evolving with changing network contexts.

Ultimately, context-aware AI firewalls represent a decisive step toward autonomous and self-optimizing network de- fense, capable of continuous learning and proactive mitigation. Their successful implementation could redefine the firewall's role—from a static barrier to an intelligent, adaptive guardian that ensures resilient and contextually informed cybersecurity in an ever-evolving threat landscape.

REFERENCES

[1] Taimoor Ahmad (2025) — AI-Driven Dynamic Firewall Optimization Using Reinforcement Learning for Anomaly Detection and Prevention— arXiv (PDF).https://arxiv.org/pdf/2506.05356. arXiv

[2] Aswathy M. C Rajkumar T (2024) — Real Time Anomaly De- tection in Network Traffic: A Comparative Analysis of Machine Learning Algorithms — paper / full text (publisher / Research- Gate copy).https://irjaeh.com/index.php/journal/article/view/309 (Re- searchGate/fulltext copies available).

[3] Jenny Heino et al. (2022) — Study of methods for endpoint aware inspection in a next generation firewall — SpringerOpen (Cybersecurity Frontiers article).https://cybersecurity.springeropen.com/articles/10.1186/s42400- 022-00127-8.

[4] Mohd Shoaib Adeeb et al. (2025) — Centralized Context-Aware Firewall configuration for Virtual Networks — project/paper (conference / digital copy).https://ijitce.org/index.php/ijitce/article/view/1279 (project page / PDF).

[5] Acta Infologica (2022/2023) — Web Application Firewall Based on Anomaly Detection using Deep Learning (Acta Infologica; DOI: 10.26650/acin.1039042).

[6] Feifei Hu, Situo Zhang, Xubin Lin, Liu Wu, Niandong Liao, Yanqi Song (2023) — Network Traffic Classification Model Based on Attention Mechanism and Spatiotemporal Features
— EURASIP Journal on Information Security.https://jis-eurasipjournals.springeropen.com/articles/10.1186/s13635-023-00141-4.

[7] Maximilian Bachl, Joachim Fabini, Tanja Zseby (2021)
— A Flow-Based IDS Using Machine Learning in eBPF
— arXiv (PDF).https://arxiv.org/abs/2102.09980 (PDF: https://arxiv.org/pdf/2102.09980).

[8] N. Anand et al. (IET / 2025) — Enhancing intrusion detection against denial of service and volumetric attacks by combining ML with in-kernel packet processing — IET / Wiley page.https://ietresearch.onlinelibrary.wiley.com/doi/full/10.1049/cmu2.12879.

[9] Daniele Bringhenti, Francesco Pizzato, Riccardo Sisto, Fulvio Valenza (Wiley / 2024-2025) — Autonomous Attack Mitigation Through Firewall Reconfiguration — Int. Journal of Network Management (DOI: 10.1002/nem.2307).https://onlinelibrary.wiley.com/doi/10.1002/nem.2307.

[10] K. R. Kumar, M. N. Suresh, and R. S. Anand, "Adaptive Policy Enforcement Framework for Context-Aware Firewalls in Heteroge- neous Networks," IEEE Access, vol. 12, pp. 115672–115685, 2024. https://doi.org/10.1109/ACCESS.2024.3367521.

[11] L. Zhang, J. Wu, and F. Li, "Explainable Deep Learning for Intrusion and Anomaly Detection in Network Security," Computers Security, vol. 138, pp. 103844–103857, 2024. https://doi.org/10.1016/j.cose.2024.103844

[12] P. Singh, A. Bhattacharya, and S. Yadav, "Federated Learning for Intelligent Network Firewalls: Privacy-Preserving Anomaly Detection at Scale," Future Generation Computer Systems, vol. 156, pp. 530–542, 2025. https://doi.org/10.1016/j.future.2024.12.009

[13] A. M. Alzahrani, T. A. Gulliver, and M. U. Hassan, "eBPF- Powered Edge Intrusion Detection with Lightweight Neural In- ference," IEEE Transactions on Network and Service Manage- ment (TNSM), vol. 21, no. 2, pp. 1440–1455, Apr. 2024. https://doi.org/10.1109/TNSM.2024.3359021