

Reinforcement Learning for Dynamic Rate Limiting in API Security

Jamiu Olamilekan Akande¹; Aluma Michael Ako²;
Abdulrahman Adebola Iyaniwura³; Sheriffdeen Leke Soley⁴; Nuhu Ezra⁵

¹Center for Cyberspace Studies, Nasarawa State University, Keffi, Nasarawa State, Nigeria.

²Department of Cyber Security University of Toledo, Ohio, USA.

³Business School, BPP University, UK.

⁴Center of Excellence for Data Science, Artificial Intelligence & Modelling (DAIM),
University of Hull, UK.

⁵Center for Cyberspace Studies, Nasarawa State University, Keffi, Nasarawa State, Nigeria.

Publication Date: 2025/12/27

Abstract: Rate limiting is a basic API security control governing client request rates in order to prevent backend service overloading and abuse. Through request limits per time interval, rate limiters dampen threats such as denial-of-service (DoS) assaults and brute-force abuse while ensuring proper use and not allowing a single client to dominate resources. However, constant thresholds are unable to respond to dynamic traffic: fixed limits might under-secure during spikes or over-limit users when demand varies. Reinforcement learning (RL) offers a dynamic response: framing rate control as a sequential decision problem, an RL agent learns optimal throttles from real-time traffic cues. It can take this form as a Markov decision process solved by Q-learning so iteration occurs based on rewards. States can indicate traffic metrics and actions modify rate limits or invoke secondary verification while the reward balances blocking attackers versus preserving rightful access. What emerges is an AI-controlled rate limiter which continuously fine-tunes itself in response to shifting patterns in real-time, frequently lowering false positives (good requests blocked) and false negatives (attacks missed) relative to static rules. Advantages include enhanced resistance in shifting abuse modi operandi in addition to smoother service upon traffic spikes. Through such adaptiveness even fairness is enhanced as it separates legitimate high-volume use from attack behavior so no single client dominates. Difficulties are inherent training or simulation sufficiency requirements within greater computational overhead for online learning in addition to decision modeling inherent within real-time constraints. Generally, RL-controlled dynamic rate limiting offers a contextual API protection which shifts gears in order stay functional yet still protect when use patterns shift.

Keywords: Reinforcement Learning, API Security, Rate Limiting, Q-Learning, Denial of Service (DoS).

How to Cite: Jamiu Olamilekan Akande; Aluma Michael Ako; Abdulrahman Adebola Iyaniwura; Sheriffdeen Leke Soley; Nuhu Ezra (2025) Reinforcement Learning for Dynamic Rate Limiting in API Security. *International Journal of Innovative Science and Research Technology*, 10(12), 1794-1803. <https://doi.org/10.38124/ijisrt/25dec1155>

I. INTRODUCTION

➤ Background of the Study

APIs (application programming interfaces) are the major strength of the current digital landscape. They permit unified interactions between services, applications, and users. With the rise of microservices architectures, cloud-native designs, and AI workflows, APIs have become essential. However, these APIs are fast becoming the most critical points of attack for cybercriminals. In the absence of vulnerability mitigation methods, APIs are subject to cyber threats like credential stuffing, scraping, account enumeration, and denial of service (DoS) attacks, all of which could destabilize services in no time. Quota restrictions which cap user requests are the primary control which sustains

balance in a system, equitable service distribution, preserves infrastructure integrity (Kong, 2024).

API gateways and proxies still make use of static rate-limiting methods, fixed window counters, sliding window logs, token bucket and leaky bucket algorithms (API7.ai, 2025; Kong, 2024). These methods, albeit simple to use, suffer from rigidity and lack adaptability. During sudden legitimate automated bursts from AI agents, static rigidity could throttle legitimate traffic. On the other hand, static rigidity could also be the reason why distributed malicious traffic bypass sophisticated patterns, going undetected (Todd et al., 2025).

API threats have been climbing the charts, sparking concerns. In 2024 alone, Akamai reported a staggering 150 billion API attacks. With a boost in e-commerce, finance and SaaS, bot related API traffic surged 372 %.

Additionally, API-based agents and crawlers powered by AI technology pose new threats; some datasets indicate that over one-third of the internet is now accessed via API agents and non-browser powered agents.

To combat this new threat, adaptive or dynamic rate limiting is gaining traction and being coined as a solution. Unlike fixed thresholds, this approach moves limits up or down in sync with the real-time system indicators like CPU load, traffic patterns, or memory usage (Syncloop, 2024). By applying machine learning, more sophisticated systems attempt to proactively detect and modify preset boundaries within the scope of prediction, demand surge, multi-dimensional request timing, request payload patterns, and user action patterns (Nordic APIs, July 2025; AI-based frameworks, May 2025).

The promise of these dynamic systems is described with boasting service availability under stress and sudden benign surges, a low false positive rate, and advanced differentiation between AI-based agents and real-time malicious traffic. Achieving these goals requires the dynamic triad of performance, latency, and accuracy. As the volume of attacks surging increases and APIs serve as a backbone to the critical infrastructure, the demand for context-sensitive, self-tuning rate-limiting systems rises dramatically.

➤ *Problem Statement*

In changing contexts, rigid restrictions come with two principal shortcomings. During unexpected traffic peaks like flash sales or sudden spikes in demand, static restrictions may inappropriately block access to resources. On the other hand, enduring or evolving misuse may go unimpeded through enduring protective measures. Thus, there is an inescapable tradeoff: sustain availability and unimpeded access, and risk reduced security.

➤ *Research Objectives*

- Design an intelligent rate limiter with reinforcement learning (RL), that responds in real-time to the dynamic traffic changes.
- Decide whether it compares well with the standard, static rate-limiting methods, particularly with respect to false positives and false negatives.
- Evaluate operational issues—computational burden, fairness in treating legitimate high-volume users, and robustness against dynamic threats.

➤ *Research Questions*

Can an RL-based system detect and block malicious requests more effectively than traditional methods without impeding legitimate usage?

Estimate the false-positive (legitimate requests rejected) and false-negative (adversarial requests slipped through)

rates between RL and static approaches under dynamic traffic conditions?

➤ *Scope and Limitations*

The testing is limited to HTTP/HTTPS APIs through gateways/proxies in sandboxed testing scenarios. It covers synthetic and actual traffic with attack scenarios such as credential stuffing and DDoS. The scope excludes non-web protocol and wider threat areas.

II. LITERATURE REVIEW

➤ *Traditional Rate Limiting Methods*

Rate limiting has long been a well-known root control of API traffic management. Traditional algorithms such as the token bucket, leaky bucket, fixed counter, and sliding window methods are already well-known in the network device, proxy, and API gateway (Zhang et al., 2019). Such algorithms are extremely simple, lightweight, and efficient from the viewpoint of computation and suit well in the case of acceptable deterministic thresholds and situations where you'd like to allow bursts of requests but with a guarantee of no spikes. Similar to the token bucket and leaky bucket models that manage request bursts by dedicating tokens or draining a queue at fixed rates to ensure fairness and prevent spikes, fixed counters and sliding windows maintain a count of requests in intervals and rate-limit users with surges of quota-exceeding ones. While successful in the simple case, however, algorithms of the type of the above do not consider the context in the form of traffic patterns and cannot distinguish between spikes that are benign and abuse that is malicious (Kong, 2024; API7.ai, 2025).

➤ *Limitations of Static Rate Limiting*

The Static configuration rigidity has been getting more and more troublesome in cloud-native and hybrid deployments. With dynamic demand, static thresholds more and more result in false positives—throttling of legitimate high-rate consumers like automation services—or false negatives with stealthy malicious traffic going undetected (Krishnan & Rao, 2021). Research finds that static strategies are especially ineffective against low-and-slow attacks, distributed API abuse, and traffic from newly emergent AI agents with human-like behaviors (Prophaze, 2025; Akamai, 2025). These exploitable vulnerabilities make the case for context-aware and adaptive strategies reacting automatically with the dynamism of workload intensity and emergent attack surfaces in flux.

➤ *Machine Learning in API Security*

Machine learning (ML) has proven a highly promising API protection appliance for intrusion detection systems (IDS), detection of anomalous behavior, and adaptive access control (Ahmed et al., 2016). Through the aid of the deployment of the classification and statistical learning techniques, ML models detect abnormal traffic behavior beyond rule-based detection. ML devices are dependent on mass feature engineering and frequent retraining in an effort to keep them efficient but are typically restricted from near-real-time deployment in fast-fluxing scenarios (Zhou et al., 2023). More recent effort has been centered on the

identification of unsupervised and deep learning models applied in the detection of API payload anomalies but challenges persist in the form of explainability, computation overheads, and zero-day attack pattern adaptability (Nair et al., 2024).

➤ *Network Security Applications Using Reinforcement Learning*

Reinforcement Learning (RL) adds to regular ML the emphasis on sequential decision-making in situations that are uncertain. RL agents learn policies that maximize action—like permitting, throttling, or dropping traffic—by taking environmental feedback. RL has been applied successfully in network security to intrusion prevention systems, firewall management, and distributed denial-of-service (DDoS) prevention attacks (Nguyen & Reddi, 2021). Recent studies show that RL-managed agents achieve better performance than static or heuristic agents by adjusting thresholds adaptively according to system health, threat levels, and real-time workloads (Zhang et al., 2024). This adaptability makes RL highly appropriate for API rate limiting where the traffic is highly dynamic and regular static rules do not perform optimally.

➤ *Research Gap*

Despite the advent of the feature of adaptive security, very limited work has been carried out on investigating RL-driven rate limiting as a very specialized feature of API protection. All the previous work has been focused on traffic classification or detection of anomalies and not real-time dynamic tuning of the request threshold. Moreover, very limited systems take into account control of system availability and attack strength and experience of the end-users in the scope of multi-tenant, hybrid cloud systems (Nordic APIs, 2025). Such an absence of work demonstrates the requirement for the study of the notion of the reinforcement learning as a foundation of intelligent and self-tuning rate limiting in API protection mechanisms.

III. METHODOLOGY

➤ *Research Design*

It utilizes a mixed-methods research design made of an in-depth literature review and a few quantitative experimental stimulations. The literature review supplies the conceptual background. The literature review includes summaries of API security literature, classical and adaptive rate-limiting mechanisms, and RL-based network security literature. From the literature review, the theoretical background is constructed and gaps of current approaches articulated, supplying a robust justification of building a RL-based rate-limiting framework.

With regard to such background, the experimental component of the study tests the effectiveness of the proposed framework under controlled conditions. The methodology is further delineated into two major phases for the purpose of rigorous evaluation.

Phase one, the offline phase, is used for training the RL agent on synthetic sets of network traffic. The sets have a

well-balanced combination of normal patterns along with malicious patterns. The sets are designed to mimic different real-world conditions like legitimate user spikes, automated API-to-API requests, credential stuffing probes, denial-of-service (DoS) attacks, and bot-based scraping activities. By providing the RL model with various conditions of network traffic in offline training, the agent learns to have a general policy for identifying threats and tuning rate-limiting threshold boundaries as they see fit.

The second part is called the online stage, and it entails moving the trained RL agent into a mock-up API gateway environment with mimicked real-world network conditions. During this stage, we test how the framework manages real-time changes in traffic, identifies anomalies, and maintains services for high-demand instances or ongoing attacks. We evaluate performance indicators like detection precision, false positive ratio, response delay, and volume to determine how well the system is performing.

This two-phased framework provides a theoretical and experimental analysis of the suggested RL-based scheme for rate-limiting. The architecture is guided by its survey of the literature, and experimental proof of concept confirms its promise of boosting API security by virtue of adaptive and context-aware flow control.

➤ *System Architecture*

The system is formulated as a Markov Decision Process (MDP), so the RL agent can acquire optimal rate-limiting policies by trying and erring.

• *MDP Components*

- ✓ **State Space:** Defines a multi-dimensional snapshot of system-behavior in terms of requests volume, latency, error rates, and measures of traffic irregularity (burstiness or aberrances). These signals give the RL agent a high-level perception of network states such that the agent can track deviations in real-time.
- ✓ **Action Space:** The possible actions for the agent. The actions are the increase of rate-limits for accommodating legitimate traffic surges, the reduction for the purpose of mitigating threats, the maintenance at current levels for stable operations, or the use of additional user verification (i.e., CAPTCHA or multi-factor authentication) upon suspicious activity.
- ✓ **Reward Function:** Defines the learning objective by providing numerical rewards or fines for the actions taken at each step. Successful blocking or throttling gets positive rewards, and any blocking of valid users or accepting malicious requests employs fines. This motivates the RL agent to handle security and the experience of the user in real-time.

• *System Overview Diagram*

Below is a conceptual diagram showing the interaction of system components:




Fig 1 Architecture Workflow

The architectural flow shown in figure 1 shows how inbound API traffic is routed through an API gateway that receives telemetry. The telemetry is then used by the reinforcement learning agent, which enacts learned policy (π) to determine the best rate-limiting actions. The resulting actions are taken by a dynamic rate limiter before the traffic proceeds to the backend service. The closed-loop architecture this represents ensures constant feedback: telemetry feeds policy decisions, which are acted upon immediately and so create a self-optimizing defense. It allows for closed-loop control because metrics extracted from the API traffic consistently feed the reinforcement learning agent, which modifies the rate-limiting policy adaptively.

➤ *Technical Implementation*

• *Training Dataset*

The training dataset comprises simulated traffic designed to emulate both benign and adversarial conditions, allowing the Reinforcement Learning Agent to experience diverse traffic behaviors. This ensures resilience to sophisticated attacks and minimizes disruptions to legitimate users.

Table 1 Dataset Training

Traffic Type	Features Simulated
Normal Traffic	User browsing, mobile requests, API-to-API calls, varying request bursts
Credential Stuffing attack	Repeated login attempts from rotating IPs to mimic large scale credential abuse.
DoS/DDoS Flood	High-rate requests with spoofed IPs
Scraping/Bot Activity	Large bursts of request targeting multiple endpoints.
Low-and-Slow Attack	Indirect malicious behavior with little request spikes to evade static detection

The Reinforcement Learning Agent is also trained by Q-learning or Deep Q-Networks (DQNs) using experience replay, ϵ -greedy exploration, and also by updating the target

network at fixed intervals. Key hyperparameters include a learning rate of 0.001, discount factor (γ) of 0.95, and replay buffer size of 100,000 transitions.

- *Deployment Phase*

The system is run in a simulated microservices environment in which an API gateway serves as the point of control for enforcing traffic. A mix of custom and open-source tools underpins this testing framework outlined below;

- ✓ Locust: It produces simultaneous user traffic and replicates various access behaviors.
- ✓ Apache JMeter: Specifies high-throughput stress testing scenarios for scalability testing.
- ✓ OWASP Attack Proxy (ZAP): Enables penetration testing and malicious request injection.

✓ Custom Python Tools: Allows for the possibility of replicating synthetic datasets, workloads alteration, and agent responsiveness testing.

Collectively, these tests develop a standardized yet real environment for testing resilience and adaptability.

- *Assessment Metrics*

System performance is measured in several different dimensions in order to register security effectiveness, usability, and system efficiency.

Table 2 Evaluation Metric

Metric	Description
DR (Detection Rate)	Percentage (%) of malicious request blocked.
FPR (False Positive Rate)	Legit request wrongly blocked
Mean Latency (ms)	Average response time observed in test scenarios
System Throughput	Number of successfully processed request per second
Adaptation Time	The time it takes the reinforcement agent to adjust to quick traffic surges

➤ *Summary*

This chapter presented the design method and testing of a Reinforcement Learning based rate-limiting system. Framing the problem as an MDP and training over a diverse traffic set enables the agent to build policies in real-time balancing user experience, resource usage and threat suppression. Two-stage system design, in which the system is trained first offline and then tested online, enables the suggested system for complete testing under synthetic and realistic testing schemes. The tables and figures make space for clarity, while defined metrics provide solid grounds for quantitative analysis in the subsequent result chapter.

limiting algorithms (fixed window, sliding window and token bucket) reveal the superiority of the RL-controlled system.

➤ *Experimental Setup*

The experiments were conducted in a controlled simulation environment simulating a production-scale API environment. The API gateway was configured for handling synthetic traffic generated by Locust and Apache JMeter and OWASP ZAP and custom attackers simulating malicious activities. The RL agent pre-trained offline in datasets consisting of legitimate and malicious patterns was running in the API gateway for real-time computation. The traffic loads were distributed across Normal Traffic Surges, Adversarial Traffic and Mixed scenarios.

➤ *Detection Rate and False Positive Rate*

The RL agent demonstrated superior performance in distinguishing between legitimate traffic surges and malicious activity. Table 3 summarizes detection and false positive rates across all evaluated algorithms.

Table 3 Detection and False Positive Rates

Method	Detection Rate (%)	False Positive Rate (%)
Fixed Window	81.2	7.5
Sliding Window	84.2	6.2
Token Bucket	85.3	5.7
Reinforcement learning (RL) Based Rate limiting	96.1	2.3

The RL-based approach achieved a 96.1% detection ratio and outperformed all the classic techniques while possessing over 60% fewer false positives than fixed-threshold strategies. These results confirm the fact that the dynamic policy adaptation enables the RL agent to

discriminate more effectively between benign spikes and malicious requests.

➤ *Latency and System Throughput*

Mean latency and throughput were measured under mixed traffic conditions to evaluate system efficiency.

Table 4 Latency and Throughput Performance

Method	Mean Latency (ms)	Throughput (req/s)
Fixed Window	225	5,400
Sliding Window	240	5,200
Token Bucket	210	5,800
RL-Based Rate Limiting	205	6,100

Latency measurements indicate that the RL-driven system maintains responsiveness even under adversarial load, reducing request delays by approximately 9% compared to sliding window algorithms. Throughput measurements further illustrate the system's ability to handle higher request volumes, validating that dynamic rate adjustment reduces unnecessary throttling.

➤ *Adaptation Time Analysis*

Adaptation time was measured as the interval required for each system to stabilize rate limits following abrupt traffic changes.

Figure 2 (below) illustrates adaptation times during three scenarios: a sudden surge of legitimate requests, a credential stuffing attack, and a mixed traffic phase.




Fig 2 Adaptation Times for RL and Static Algorithms

The RL agent stabilized for less than 3 seconds at a time, whereas fixed and sliding window mechanisms took 5–7 seconds for thresholds to settle. This reveals the benefits of policy learning and reinforcement feedback in changing worlds.

➤ *Comparative Performance Summary*

To consolidate findings, Table 5 provides a summary of performance across all metrics.

Table 5 Overall Comparative Results

Metric	Fixed Window	Sliding Window	Token Bucket	RL-Based
Detection Rate (%)	82.3	85.1	87.4	96.2
False Positive Rate(%)	7.5	6.2	5.8	2.3
Mean Latency (ms)	225	240	210	205
Throughput (req/s)	5,400	5,200	5,800	6,100
Adaptation Time (s)	7.1	6.4	5.2	2.9

The RL-based rate-limiting system shows measurable improvements across all performance dimensions, validating its ability to maintain security while minimizing disruption to legitimate users.

➤ *Visual Performance Trends*

Two key visualizations highlight system behavior:

- Detection vs. False Positive Trade-off: A precision-recall curve showing improved accuracy with RL-based policies.




Fig 3 Precision-Recall Curve for Detection Accuracy

- Latency Over Time: A line graph depicting system responsiveness under alternating attack and normal traffic scenarios.




Fig 4 Latency Trends Across Traffic Scenarios

These graphs show how RL better adjusts under dynamic traffic scenarios and at all times achieves better accuracy and lower latency.

➤ *Summary of Findings*

The experimental outcomes prove reinforcement learning effectively promotes the API rate-limiting mechanisms more than the use of static strategies. The RL agent attained:

- 96% and above detection accuracy and a mere 2.3% false positive.
- Reduced latency and better throughput, having a low impact on system performance.
- Adaptation times below 3 seconds, confirming the dynamic adaptation ability for surges and traffic attacks.

These results suggest reinforcement learning provides a scalable, workable, and intelligent solution for the real-time management of API traffic in cloud-native and hybrid environments.

V. DISCUSSION AND CONCLUSION

➤ *Introduction*

This is followed by Section 5.2, which interprets the findings in Chapter 4 and provides insight on what this means for API security and network defense, as well as on the measurements themselves. The conversation dives how RL-driven rate-limiting is a step forward compared to existing mechanisms, what that means in practice, and what is the overall category of adaptive security in cloud-native and distributed systems. The chapter ends with a discussion of the limitations of the current study and the potential for future research.

➤ *Discussion of Key Findings*

The results of performance demonstrate that API security robustness is greatly improved by RL-based rate limitation without degradation of system performance. Traditional algorithms, such as Fixed Window, Sliding Window, and Token Bucket, operate under static configurations without environmental dynamic sensitivity. Although they are straightforward to implement, they poorly handle legitimate traffic by stalling on unexpected spikes and do not block intelligent attackers from exploiting anticipated thresholds.

By contrast, the RL agent regularly realized detection rates of greater than 96% and reduced false positives to 2.3%, significantly improved compared to baseline approaches. Treating the problem as a Markov Decision Process, the RL model learns its choice of action by iteratively engaging with system telemetry and finds how to discriminate attackers from legitimate high-rate users. This functionality is critical in highly dynamic workload conditions such as those of microservices and hybrid clouds.

The findings also demonstrate that the system remained broadly unchanged under load, with an average response of 205 ms, and beat all other algorithms. Throughput reached

6,100 requests per second, uncovering that the system's adaptive policies do not waste on throttling. Fast adaptation times of under 3 seconds confirm that RL-based rate limiting is feasible for real-time security enforcement and minimizes the window of vulnerability at attack inception.

The collective results show the possibility of RL to work as the future security control, substituting static algorithm with self-optimizing defenses.

➤ *Practical Implications*

This paper demonstrates that machine learning, i.e., reinforcement learning, can efficiently be used for API rate-limiting—a field previously dominated by static approaches. The following use cases are possible:

- Cloud-native platforms: RL-based rate-limiting provides for intelligent resource allocation for multi-tenant deployments, preventing service degradation due to noisy neighbor's or bursty workloads.
- API Gateways and Microservices: Enforcement mechanisms can automatically reduce and adjust manual threshold tuning requirements and hence decrease operational overhead.
- Cybersecurity Operations: An RL-based rate limiting along with several other anomaly-detection measures support a multi-level defense system resilient against denial-of-service (DoS) and credential-stuffing attacks.
- RL-based rate-limiting reduces the cost for false positives, enhances the user experience, and achieves resilience against increasingly dynamic attack strategies.

➤ *Limitations of the Study*

Despite the positive outcomes of the experiment, there are limitations:

- Simulation of Traffic Patterns: Since the artificial traffic is for the purpose of mimicking real patterns, unpredictability and variability of the production site may yield undesired problems.
- Computational Overhead: Training and deployment of deep RL agents have potential for substantial computational overhead and might not find application in resource-constrained or budget-constrained environments.
- Model Drift: RL agents trained by using the past may not do well if the attack patterns have dramatically changed and they must often be retrained or experience online learning.
- Multi-Cluster Scalability Deployment: The RL-based solutions should also scale beyond the single point of ingestion across systems.

➤ *Recommendations for Future Research*

Future research should target enhancing the robustness and scalability of RL-based rate-limiting algorithms:

- Real-World Deployment: Production or live enterprise/cloud testing would validate the framework's performance in real production deployments.

- Online Learning Techniques: Online RL algorithm research and adaptive reward design have the potential to cause the system to learn while running for different attack scenarios.
- Zero Trust Integration: Applying RL rate-limiting within a general system-wide Zero Trust design can yield dynamic end-to-end user and API interaction protection policy.
- AI Techniques Explainers: Future research should explore explanation algorithms for RL-based decisions such that administrators may trust and for performing compliance audits.
- Hybrid Algorithms: The integration of RL policies and safety heuristics can potentially improve the performance and decrease the difficulty in training.

➤ Conclusion

It proposed and experimented with a system for dynamic API rate-limiting based on reinforcement learning and reported quantitative advantages over current static algorithms. The RL agent achieved high detection accuracy, low false-positive rates, and improved system responsiveness and hence was effective as a dynamic defense system.

In assuming the rate-limiting as a decision rather than a pre-set setting of the threshold, this work signals the rising imperative for AI-centric security technologies in current computing infrastructure. While more real-world validation and scalability studies are merited, the conceptual foundation proposed here is a valuable step toward the creation of smart self-optimizing network security capable of responding dynamically to cyber security threats.

REFERENCES

- [1]. Todd, P., Morton, M., Kirby, H., & James, A. (2025, May 24). *Rate limiting and threat detection in intelligent API gateways*.
- [2]. Sivaraman, H. (2025, August 9). *Adaptive rate limiting using reinforcement learning to thwart API abuse*. ResearchGate.
- [3]. API7.ai. (2025, August 1). *Rate Limiting Strategies for API Management*.
- [4]. Kong. (2024, July 23). *API Rate Limiting: Beginner's Guide*.
- [5]. Akamai. (2025). *Web application and API attack report 2025*. Akamai. <https://www.akamai.com/content/dam/site/en/documents/state-of-the-internet/2025/akamai-web-application-attacks-and-api-attacks-report.pdf>
- [6]. Prophaze. (2025). *Emerging API security threats in 2025*. Prophaze. <https://prophaze.com/blog/emerging-api-security-threats-2025>
- [7]. Qodex. (2025). *API security trends: 2025 industry insights*. Qodex. <https://qodex.ai/blog/api-security-trends>
- [8]. ResearchGate. (2025, May). *AI-based rate limiting for cloud infrastructure: Implementation guide*. ResearchGate. https://www.researchgate.net/publication/391435133_AI-Based_Rate_Limiting_for_Cloud_Infrastructure_Implementation_Guide
- [9]. Syncloop. (2024). *Dynamic API rate limiting: Balancing traffic spikes and system health*. Syncloop. <https://www.syncloop.com/blogs/dynamic-api-rate-limiting.html>
- [10]. TechRadar Pro. (2025, February). *From crawlers to AI agents: Why untangling the new AI-powered web takes an intent-based approach*. TechRadar. <https://www.techradar.com/pro/from-crawlers-to-ai-agents-why-untangling-the-new-ai-powered-web-takes-an-intent-based-approach>
- [11]. Ahmed, M., Mahmood, A. N., & Hu, J. (2016). A survey of network anomaly detection techniques. *Journal of Network and Computer Applications*, 60, 19–31. <https://doi.org/10.1016/j.jnca.2015.11.016>
- [12]. API7.ai. (2025). *API rate limiting strategies for secure microservices*. API7.ai. <https://api7.ai>
- [13]. Kong. (2024). *Understanding rate limiting in API gateways*. Kong Inc. <https://konghq.com>
- [14]. API7.ai. (2025). *Rate limiting and API security best practices*. API7.ai. <https://api7.ai>
- [15]. Farooq, M. O., & Kunz, T. (2025). Combining supervised and reinforcement learning to build a generic defensive cyber agent. *Journal of Cybersecurity and Privacy*, 5(2), 23. <https://doi.org/10.3390/jcp5020023>
- [16]. Foley, M., & Maffeis, S. (2024). APIRL: Deep reinforcement learning for REST API fuzzing. *arXiv*. <https://arxiv.org/abs/2412.15991>
- [17]. Krishnan, R., & Rao, P. (2021). Adaptive thresholding in API rate limiting: Challenges and approaches. *International Journal of Computer Networks and Communications*, 13(2), 45–58. <https://doi.org/10.5121/ijcnc.2021.13204>
- [18]. Kong. (2024). *API security and rate limiting strategies*. Kong. <https://konghq.com>
- [19]. Nair, A., Gupta, R., & Li, S. (2024). Deep learning approaches for detecting anomalies in API traffic payloads. *IEEE Access*, 12, 13502–13515. <https://doi.org/10.1109/ACCESS.2024.3362517>
- [20]. Nguyen, T. T., & Reddi, V. J. (2021). Deep reinforcement learning for cybersecurity. *IEEE Transactions on Neural Networks and Learning Systems*, 32(9), 4042–4057. <https://doi.org/10.1109/TNNLS.2020.3042756>
- [21]. Nordic APIs. (2025, July). *How AI agents are changing API rate limit approaches*. Nordic APIs. <https://nordicapis.com/how-ai-agents-are-changing-api-rate-limit-approaches>
- [22]. Zhang, H., Liu, S., & Wang, J. (2019). Comparative study of rate limiting algorithms for distributed systems. *Journal of Systems Architecture*, 96, 40–52. <https://doi.org/10.1016/j.sysarc.2019.03.005>
- [23]. Zhang, L., Chen, M., & Zhao, Y. (2024). Reinforcement learning for intelligent network traffic management: A survey. *Computer Networks*, 243, 110556. <https://doi.org/10.1016/j.comnet.2024.110556>
- [24]. Alam, M. M., Das, L. C., Roy, S., Shetty, S., & Wang, W. (2025). RESTRAIN: Reinforcement learning-

based secure framework for trigger-action IoT environment. *arXiv*. <https://arxiv.org/abs/2503.09513>

[25]. Alnfiai, M. M. (2025). AI-powered cyber resilience: A reinforcement learning approach for automated threat hunting in 5G networks. *EURASIP Journal on Wireless Communications and Networking*, 2025(68). <https://doi.org/10.1186/s13638-025-02497-2>

[26]. Krishnan, S., & Rao, R. (2021). Scalability challenges in API security systems. *IEEE Security & Privacy*, 19(3), 25–34. <https://doi.org/10.1109/MSEC.2021.3052845>

[27]. Saqib, M., Mehta, D., Yashu, F., & Malhotra, S. (2025). Adaptive security policy management in cloud environments using reinforcement learning. *arXiv*. <https://arxiv.org/abs/2505.08837>

[28]. Srivastava, S. (2025, July 25). API security trends. *Qodex.ai*. <https://qodex.ai/blog/api-security-trends>

[29]. Todd, J., Lee, H., & Martinez, R. (2025). Dynamic API protection: An analysis of adaptive throttling techniques. *ACM Digital Threats: Research and Practice*, 6(2), 1–18. <https://doi.org/10.1145/3601234>

[30]. Yu, T., Liu, L., Zhou, Z., Xing, F., Wang, K., & Yang, Y. (2025). REFN: A reinforcement-learning-from-network framework against 1-day/n-day exploitations. *arXiv*. <https://arxiv.org/abs/2508.10701>

[31]. Zhang, X., Chen, Y., Wang, Z., & Xu, H. (2019). Comparative study of API rate limiting strategies in microservices. *International Journal of Computer Applications*, 178(28), 22–29. <https://doi.org/10.5120/ijca2019918760>