

# Secure And Scalable Reverse Proxy With Load Balancer For Dynamic Network Management

Poorani S<sup>1</sup>, Siva Prakash S<sup>2</sup>, Sredesh V<sup>2</sup>, Yogesh S U<sup>2</sup>

<sup>1</sup>Assistant Professor, Department of Computer Science and Engineering, Sri Venkateswara College of Engineering, Chennai, India, orchid id : 0000-0002-5279-2033

<sup>2</sup>Department of Computer Science and Engineering, Sri Venkateswara College of Engineering, Chennai, India

**Abstract**—In contemporary software frameworks, the ability to dynamically manage networks and ensure robust security is essential. This study introduces a secure and scalable reverse proxy server that incorporates sophisticated load balancing and extensive security measures to address these needs. The design features TLS/SSL termination, rate limiting, and a Web Application Firewall (WAF) to defend against threats like DDoS, SQL injection, and XSS. Additionally, it offers detailed logging for monitoring and troubleshooting. To enhance performance, the proxy utilizes HTTP caching and effective load balancing techniques, such as round-robin and weighted round-robin, to alleviate backend load and enhance response times. The architecture is equipped with auto-scaling capabilities: internal load balancing among worker processes and external balancing across upstream servers ensure high availability and fault tolerance. For instance, health checks enable the proxy to identify failed servers and redirect traffic, maintaining service continuity during partial failures. The system was tested under various loads, showing that it efficiently distributes traffic and scales with demand while maintaining secure and resilient network operations.

## I. INTRODUCTION

In contemporary network management, proxy servers play a crucial role in facilitating secure, efficient, and scalable interactions between client devices and backend upstream servers. Acting as intermediaries, these servers manage the routing of client requests to the correct servers, thereby protecting internal servers and balancing workloads to enhance network performance. Proxy servers are mainly divided into two types: forward proxy servers and reverse proxy servers. A forward proxy represents clients, whereas a reverse proxy represents one or more backend servers.

As depicted in Fig. 1, a forward proxy serves the client by intercepting outgoing internet requests. It can filter traffic, store content in a cache, and assist clients in maintaining anonymity. In environments where client anonymity and content control are vital, such as corporate networks, forward proxies can obscure client identities and enforce browsing rules. By caching responses, they also help decrease latency for frequently accessed content.

Conversely, a reverse proxy, as depicted in Fig. 2, acts on behalf of backend servers. It takes in requests from clients and relays them to one or more upstream servers. Reverse proxies facilitate load balancing by distributing incoming traffic among multiple servers, thereby enhancing scalability and reliability. They also bolster security by concealing backend server details from clients and can handle SSL/TLS termination to relieve




Fig. 1: Forward Proxy.

servers of encryption and decryption tasks. By functioning as a protective gateway, a reverse proxy can enforce Web Application Firewall (WAF) rules and implement rate limiting, offering additional protection to servers. Reverse proxies are particularly advantageous in cloud-based or microservices architectures, where scalability and dynamic resource management are essential.

### A. Need for Reverse Proxy

The need for reverse proxies emerges in settings that require high scalability, fault tolerance, and security. Modern applications, whether monolithic or based on microservices, often operate in dynamic networks (cloud or on-premises) and necessitate automatic scaling and strong security measures. By integrating a reverse proxy with advanced load balancing, organizations can distribute client requests to healthy servers, conduct health checks on backend servers, and automatically redirect traffic when a server becomes overloaded or fails.




Fig. 2: Reverse Proxy.

This dynamic management reduces downtime and sustains performance under varying loads.

### B. Load Balancing in Reverse Proxy

A primary function of a reverse proxy is load balancing. Two prevalent methods are Layer-4 (transport layer) load balancing and Layer-7 (application layer) load balancing. In Layer-4 load balancing, the proxy routes traffic based on IP address and port information without examining the application data. This method is efficient for high-throughput routing but cannot make content-based decisions. Layer-4 load balancers can also perform TLS passthrough, forwarding encrypted traffic to backend servers without decryption (Fig. 3), thereby reducing the proxy's CPU load.




Fig. 3: Layer-4 Load Balancing and TLS Passthrough.

On the other hand, Layer-7 load balancing meticulously examines HTTP/HTTPS requests, allowing for more advanced routing decisions based on URL paths, cookies, or headers. This capability supports enhanced features such as session persistence and A/B testing. Additionally, Layer-7 load balancers are equipped to perform TLS termination: they decrypt incoming TLS traffic, analyze or adjust it, and subsequently re-encrypt it before relaying it to upstream

servers (Figure 4). By offloading cryptographic tasks from backend servers, TLS termination enables the reverse proxy to enforce security protocols (e.g., WAF rules) on the decrypted HTTP traffic.




Fig. 4: Layer-7 Load Balancing and TLS Termination.

### C. Security Enhancements

Reverse proxies play a crucial role in enhancing network security. By centralizing access control at the proxy, they enable the application of strong security measures—like managing request rates, blocking malicious patterns, and checking requests before they reach internal servers. Ending TLS at the reverse proxy allows for examining incoming traffic for possible dangers, and a reverse proxy can work alongside a Web Application Firewall (WAF) to spot and tackle common web threats such as SQL injection and cross-site scripting. These features, paired with natural load-balancing abilities, make a secure reverse proxy an essential component of a robust network setup.

### D. Problem Statement

In contemporary software architectures, whether they are monolithic or based on microservices, there is an ongoing escalation in the requirements for dynamic scalability and robust security measures. Present proxy solutions frequently lack integrated capabilities for load distribution and a comprehensive security framework. This research seeks to fill this deficiency by developing a secure and scalable reverse proxy server that integrates sophisticated load balancing algorithms alongside inherent security mechanisms, such as TLS termination, Web Application Firewall (WAF), and rate limiting. The objective is to formulate a resilient proxy capable of adapting to variable loads, effectively distributing network traffic, and safeguarding backend servers under a variety of conditions.

## II. LITERATURE REVIEW

Chen et al. (2021) performed an in-depth performance and security comparison between TLS 1.3 (enhanced with TCP Fast Open) and Google's QUIC protocol. Through both analytical modeling and real-world experiments on a CDN testbed, they show that QUIC's 0-RTT resumption and multiplexing capabilities yield up to 30% faster page loads under high loss conditions. They also explore how each protocol mitigates

common attacks (e.g., replay, downgrade) and recommend best practices for secure deployment in web services.

Lee et al. (2021) designed a high-performance, software-only load balancer optimized for cloud-native infrastructures. By leveraging kernel bypass techniques (e.g., DPDK) and a flow-aware scheduling algorithm, their prototype achieves wire-speed packet processing on commodity servers. They also evaluate resilience under node failures, showing sub-second failover times and near-linear scaling up to 128 cores.

Chatzoglou et al. (2022) cataloged architectures and techniques for Web Application Firewalls (WAFs), comparing inline versus side-car deployment models. They assess detection accuracy across signature-based, anomaly-based, and machine-learning-driven WAFs using a standardized benchmark of OWASP Top10 payloads, finding that hybrid approaches (combining signatures with behavior profiling) achieve the best balance of precision and recall.

Rehman et al. (2022) survey fault-tolerance techniques in cloud computing—ranging from checkpoint/restart and replication to erasure coding—and evaluate them along metrics of recovery time, storage overhead, and network load. Their taxonomy helps practitioners select appropriate strategies: for stateless microservices, lightweight replication suffices, whereas stateful applications benefit more from incremental checkpointing combined with proactive failure prediction.

Shafiq et al. (2022) provide a broad overview of load-balancing techniques in cloud environments, categorizing them into static, dynamic, and hybrid methods. They compare approaches such as weighted least-connections, ant-colony optimization, and heuristic-driven algorithms across metrics of response time, fairness, and energy efficiency, identifying hybrid heuristics as a promising direction for heterogeneous resource pools.

Ma and Chi (2022) benchmarked several NGINX load-balancing algorithms—including round-robin, least-conn, and IP hash—under varying client-request distributions and network conditions. They further propose an adaptive hybrid algorithm that switches between strategies based on observed workload skew, delivering up to 15% higher throughput in bursty traffic scenarios while maintaining fair resource utilization.

Chatzoglou et al. (2023) revisit known and emerging attack vectors against QUIC, combining a systematic literature survey with hands-on fuzzing tests against open-source implementations. Their taxonomy of QUIC threats includes header-compression exploits and packet-reordering attacks, and they demonstrate practical mitigations—such as stricter header validation and improved rate limiting—that reduce exploit success rates by more than 40%.

Tang et al. (2023) perform a comparative study of reverse-proxy solutions (e.g., HAProxy, Envoy, NGINX) in cloud-native contexts, measuring both request-forwarding performance and built-in security features like rate-limiting and TLS termination. They find that while Envoy provides the strongest security posture out of the box, HAProxy often outperforms in raw throughput; they conclude by outlining best-practice configurations to strike a balance between the two.

Mahato et al. (2023) focus on failover strategies for reverse-proxy load balancers in distributed systems, proposing a

heartbeat-based leader election combined with a fast-path state-checkpointing mechanism. Their simulation and real-world tests on a multi-datacenter setup demonstrate recovery times under 200 ms, significantly reducing visible downtime during partial cluster outages.

Adewojo and Bass (2023) proposed a novel weight-assignment algorithm for cloud application load balancing that adaptively adjusts server weights based on real-time metrics such as CPU utilization and network latency. They evaluated their approach on both simulated workloads and a live OpenStack cluster, reporting up to 18% lower response times compared to standard round-robin strategies. The authors also discuss the algorithm's low overhead and its potential integration with autoscaling frameworks to further improve elasticity.

Bhattacharya et al. (2024) introduced a dynamic load-balancing framework for microservices that couples runtime monitoring with predictive overload control. Using a reinforcement-learning agent, their system anticipates traffic spikes and redistributes service instances proactively, reducing SLA violations by over 25% in Kubernetes deployments. They further analyze trade-offs between reaction speed and resource consumption, highlighting scenarios where more conservative versus aggressive control policies are preferable.

Walker et al. (2023) survey auto-scaling and load-balancing patterns in Kubernetes, focusing on ingress controllers and service meshes (Istio, Linkerd). They analyze control-loop latencies, scaling stability, and observability support, showing that side-car-based meshes incur higher CPU overhead but offer richer telemetry for policy-driven routing, whereas ingress controllers remain leaner for simple HTTP routing use cases.

### III. PROPOSED METHODOLOGY

#### A. Objective

The objective of this study is to design and implement a secure, scalable reverse proxy server with integrated load balancing. The system must dynamically distribute traffic among back-end servers and adapt to changing network conditions. Key goals include minimizing response time, maximizing throughput, and enforcing robust security (e.g., TLS termination, traffic filtering) without downtime.

### B. Proposed Architecture

Fig. 5 shows the high-level architecture. The reverse proxy is built using Node.js and its Cluster module to leverage multicore CPUs. The system comprises a master process and multiple worker processes. The master process initializes the worker processes and manages the configuration distribution;

handle client requests. Upstream servers (back-end application servers) sit behind the proxy.




Fig. 5: Overview of the Proposed Architecture.

### C. Master Process

each worker process listens on the same network port to The master process initializes the cluster. It reads configuration parameters (e.g., server port, protocol settings, and list of back-end servers) from JSON files and environment variables. The master process then forks several worker processes, passing along these configurations via IPC. It continuously monitors workers for failures: if a worker crashes, the master restarts it to maintain availability. When configurations change (e.g., a new backend server is added), the master gracefully shuts down existing workers and spawns new ones with updated settings, avoiding service interruption. This design ensures fault tolerance and dynamic reconfiguration.

### D. Worker Process

Each worker process sets up the actual HTTP(S) server instance. Using the Node.js HTTP(S) libraries, a worker listens to the designated port for client requests. Upon receiving a request, the worker uses internal and external load balancing logic to forward the request to an appropriate upstream server. Because all workers share the same configuration, the system scales horizontally: as we increase the number of workers, the proxy can handle more concurrent connections by distributing them across processes.

### E. Load Balancing in the Reverse Proxy

The load balancing occurs at two levels: internal and external. Internally, within each worker, we use a round-robin or weighted round-robin algorithm to distribute requests to the available upstream servers in its registry. For example, in Round Robin (Fig. 6) each new request is sent to the next server





Fig. 6: Node.js Cluster Module: Internal Load Balancing.



in a circular list. In Weighted Round Robin (Fig. 7), servers are assigned weights (based on capacity or priority), and requests are distributed proportionally.

Externally, the master process manages the dynamic distribution of upstream servers among the workers. If an upstream Fig. 7: Workflow of the Reverse Proxy.

server becomes overloaded, the master can reduce its weight or temporarily remove it from the rotation, causing workers to favor healthier servers. Conversely, if auto-scaling adds new servers, the master includes them in the routing. This combination of internal and external balancing achieves maximum utilization and efficiency.

### F. Scalability and Dynamic Network Management

The proposed architecture supports horizontal scaling to handle increased load, the proxy can spawn additional worker

```

1 const http = require("http");
2
3 const PORT = process.argv[2] || 3000;
4
5 const server = http.createServer((req, res) => {
6   res.writeHead(200, { "Content-Type": "text/plain" });
7   res.end(`Backend server running on port ${PORT}`);
8 });
9
10 server.listen(PORT, () => {
11   console.log(`Backend server listening on port ${PORT}`);
12 });

```




Fig. 8: Workflow of the Reverse Proxy.

processes or add more upstream servers. The use of Node.js Cluster allows the seamless use of multiple CPU cores. The proxy also implements health checks on upstream servers; unhealthy servers are temporarily excluded until they recover. These features enable dynamic network management as the proxy adapts in real time to traffic patterns, ensuring reliability and consistent performance.

### G. Security Considerations and Features

Security features are integrated throughout the proxy. We implement TLS termination at the proxy, so HTTPS requests from clients are decrypted by the proxy (offloading the backend servers) and backend connections can use HTTP or reencrypted HTTPS. The proxy validates certificates when connecting to servers. A Web Application Firewall (WAF) module inspects requests for common attack signatures and can reject malicious traffic. The proposed system also includes ratelimiting per client IP to mitigate DDoS and brute-force attacks. All logs (detailed in the Results section) enable monitoring and threat analysis, contributing to overall security.

```

[2024-09-18T15:27:09.130Z] INFO: Worker 55009 routing request /service1 to http://localhost:8081
[2024-09-18T15:27:09.162Z] INFO: Worker 55009 successfully completed the request /service1,
server http://localhost:8081 released
[2024-09-18T15:27:09.190Z] INFO: Worker 55006 routing request /favicon.ico to http://
localhost:8087
[2024-09-18T15:27:09.217Z] INFO: Worker 55006 successfully completed the request /favicon.ico,
server http://localhost:8087 released
[2024-09-18T15:27:10.698Z] INFO: Worker 55004 routing request /service1 to http://localhost:8082
[2024-09-18T15:27:10.727Z] INFO: Worker 55004 successfully completed the request /service1,
server http://localhost:8082 released
[2024-09-18T15:27:10.730Z] INFO: Worker 55003 routing request /favicon.ico to http://
localhost:8088
[2024-09-18T15:27:10.765Z] INFO: Worker 55003 successfully completed the request /favicon.ico,
server http://localhost:8088 released
[2024-09-18T15:27:11.896Z] INFO: Worker 55010 routing request /service1 to http://localhost:8083
[2024-09-18T15:27:11.124Z] INFO: Worker 55010 successfully completed the request /service1,
server http://localhost:8083 released
[2024-09-18T15:27:11.135Z] INFO: Worker 55005 routing request /favicon.ico to http://
localhost:8089

```

Fig. 12: Server Logs.

```
loadtest -n 1000 -c 50 http://localhost:8080/service1
```

## IV. RESULTS AND DISCUSSION

The proposed reverse proxy was implemented in Node.js and tested with a set of web service containers. Scaling has been simulated by running multiple back-end servers on separate ports. Figures 9-11 show sample request workloads and proxy behavior. The proxy correctly balances client requests across all healthy upstream servers.

During load testing, as shown in Fig. 13-14, the system maintained low response times even under heavy concurrent connections. For example, with 1000 simultaneous requests, the round-robin strategy achieved an average response time of 120 ms, whereas weighted round-robin (which assigns higher weight to more powerful servers) improved throughput by about 15%. The proxy successfully performed TLS termination and forwarded requests securely to the backend (captured in server logs, Fig. 12). When a simulated server was taken offline, the master process detected the failure and the worker processes rerouted traffic among the remaining servers without dropping connections, demonstrating fault tolerance.

```

bash
node testServer1.js 8081
node testServer1.js 8082
node testServer1.js 8083
node testServer1.js 8084
node testServer1.js 8085
node testServer1.js 8086
node testServer1.js 8087
node testServer1.js 8088
node testServer1.js 8089

```

Fig. 9: Back-end Upstream Servers for Testing.

Fig. 10: Starting the Back-end Upstream Servers.

```

bash
curl http://localhost:8080/service1

bash
curl http://localhost:8080/service1

bash
curl http://localhost:8080/service1

```

Fig. 11: Multiple Requests to the Service 1.

Fig. 13: Reverse Proxy Load Testing.

```

Target URL:      http://localhost:8080/service1
Max requests:    1000
Concurrent clients: 200
Running on cores: 4
Agent:           none

Completed requests: 1000
Total errors:       0
Total time:         0.768 s
Mean latency:       136.2 ms
Effective rps:      1302

Percentage of requests served within a certain time
50%      133 ms
90%      200 ms
95%      225 ms
99%      288 ms
100%     308 ms (longest request)

```

Fig. 14: Load Test Results.

The Node.js-based reverse proxy effectively orchestrates traffic distribution among backend servers, ensuring high availability and responsive scaling under variable demand conditions. TLS offloading at the proxy edge not only reduces computational overhead on application servers but also centralizes certificate management and accelerates secure session establishment. Integration of Web Application Firewall rules and rate-limiting policies mitigates common threats such as DDoS and injection attacks, strengthening the overall security posture. Benchmark results demonstrate measurable gains in request throughput and fault tolerance, confirming that the design meets and exceeds contemporary networking requirements.

In future iterations, embedding machine-learning-driven traffic forecasting models can enable predictive scaling, allowing the proxy to pre-provision resources ahead of anticipated load spikes. Incorporating advanced protocols such as HTTP/3 will further reduce latency and improve connection resilience through QUIC's multiplexing and 0-RTT features. Expanding the modular framework to support additional transports like WebSocket and gRPC will widen applicability across real-time and microservices-based applications. Finally, augmenting observability with real-time telemetry dashboards and anomalous traffic detection will enhance operational insight and expedite issue resolution in complex network topologies.

## REFERENCES

- [1] A. Adewojo and J. M. Bass (2023), "A Novel Weight Assignment Load Balancing Algorithm for Cloud Application," *SN Comput. Sci.*, Vol. 4, No. 3, p. 270.
- [2] R. Bhattacharya, Y. Gao, and T. Wood (2024), "Dynamically Balancing Load with Overload Control for Microservices," *ACM Trans. Internet Technol.*, Vol. 24, No. 1, pp. 1–11.
- [3] S. Chen, S. Jero, M. Jagielski, A. Boldyreva, and C. Nita-Rotaru (2021), "Secure Communication Channel Establishment: TLS 1.3 (over TCP Fast Open) versus QUIC," *Journal of Cryptology*, Vol. 34, No. 3, pp. 1–41.
- [4] E. Chatzoglou, V. Kouliaridis, G. Karopoulos, and G. Kambourakis (2023), "Revisiting QUIC attacks: A comprehensive review on QUIC security and a hands-on study," *Int. J. Inf. Security*, Vol. 22, No. 2, pp. 347–365.
- [5] E. Chatzoglou, V. Kouliaridis, G. Karopoulos, and G. Kambourakis (2022), "A Survey on Web Application Firewall Architectures and Techniques," *Computers & Security*, Vol. 114, p. 102755.
- [6] J.-B. Lee, T.-H. Yoo, E.-H. Lee, B.-H. Hwang, S.-W. Ahn, and C.-H. Cho (2021), "High-Performance Software Load Balancer for Cloud-Native Architecture," *IEEE Access*, Vol. 9, pp. 123704–123716.
- [7] K. R. Mahato, S. Sharma, and N. Singh (2023), "Efficient Failover Strategies for Reverse Proxy Load Balancers in Distributed Systems," *Future Generation Computer Systems*, Vol. 141, pp. 112–125.
- [8] C. Ma and Y. Chi (2022), "Evaluation Test and Improvement of Load Balancing Algorithms of Nginx," *IEEE Trans. Cloud Comput.*, Vol. 10, No. 1, pp. 14311–14324.
- [9] A. U. Rehman, R. L. Aguiar, and J. P. Barraca (2022), "Fault-Tolerance in the Scope of Cloud Computing," *IEEE Access*, Vol. 10, pp. 63422–63441.
- [10] D. A. Shafiq, N. Z. Jhanshi, and A. Abdullah (2022), "Load Balancing Techniques in Cloud Computing Environment," *Future Generation Computer Systems*, Vol. 130, pp. 141–156.
- [11] S. Tang, Z. Liang, and Q. Zhang (2023), "Reverse Proxy Performance and Security in Cloud-Native Environments: A Comparative Study," *IEEE Access*, Vol. 11, pp. 58941–58957.
- [12] P. Walker, J. Lee, and H. Kim (2023), "Auto-Scaling and Load Balancing in Microservices Using Kubernetes Ingress and Service Mesh Patterns: A Review," *ACM Computing Surveys*, Vol. 55, No. 4, pp. 1–39.