International Journal of Computational and Experimental Science and ENgineering (IJCESEN)

Vol. 11-No.4 (2025) pp. 7704-7711
http://www.ijcesen.com

ISSN: 2149-9144

Research Article

Active-Active DNS Architectures: Building Resilient Global-Scale Name Resolution Systems

Anil Puvvadi*

Independent Researcher, USA
* **Corresponding Author Email:** reachanilpuvvadi@gmail.com- **ORCID:** 0000-0002-5247-7000

Article Info:

DOI: 10.22399/ijcesen.4099
Received : 25 August 2025
Accepted : 11 October 2025

Keywords

Active-Active Architecture,
DNS Resilience,
Multi-Plane Design,
DDoS Mitigation,
Global Synchronization

Abstract:

Active-active DNS architectures are a paradigm shift in terms of the construction of resilient name resolution infrastructure that can satisfy the current needs of the Internet. Classical active/ passive type of failure mode illustrates striking weaknesses when faced with the current demands of uninterrupted availability, extreme latency, and defense against advanced attacks. The architectural development of simultaneous multi-plane operations gets rid of single failure points with systematically redundant heterogeneous technology stacks. These deployments utilize decentralized concepts, spreading zone data across stand-alone resolver planes that preserve self-operation capabilities while synchronizing over high-speed replication pipelines. Defense-in-depth techniques utilize more than one filtering layer, ranging from network-edge volumetric defense to application-layer anomaly detection, building robust protections against dynamic threat environments. Data synchronization technologies find consistency requirements and performance demands in balance through event-driven designs and cryptographic authentication protocols. Operational excellence is realized through ongoing optimization, with chaos engineering techniques confirming resilience hypotheses and remediation automation systems ensuring service continuity. The interaction of these architectural aspects allows DNS infrastructures to meet nearly perfect availability objectives while handling hundreds of billions of queries per day across distributed networks worldwide.

1. Introduction

The Domain Name System (DNS) represents the cornerstone of Internet functionality, performing the essential task of converting domain names into numerical IP addresses that network devices require for communication. Recent investigations into DNS over HTTPS (DoH) implementations reveal that encrypted DNS queries now constitute approximately 23% of total DNS traffic, fundamentally altering traditional network monitoring approaches [1]. This shift toward encrypted DNS protocols introduces significant challenges for network administrators attempting to maintain visibility into DNS resolution patterns while preserving user privacy. The emergence of DoH has necessitated the development of sophisticated detection mechanisms, with machine learning classifiers achieving detection accuracy rates between 94.7% and 98.3% across various network conditions [1].

Modern Internet infrastructure is able to respond to queries on a scale never seen before, with a single root server answering more than 940 billion queries per year, as recent traffic reports indicate. The massive expansion of interconnected gadgets, especially in Internet of Things (IoT) systems, has increased the load size on DNS architectures many times over. The SDN environment has been found to add more name queries within 340 percent over the last three years due to increased IoT deployments, particularly their use in smart homes, smart sensors, and autopilot systems that continuously need name lookups [2]. Such IoT networks present a distinct set of weaknesses, as scarcely resourced devices are frequently not protected by powerful security tools; consequently, DNS infrastructure has been pursued by intruders.
Sophisticated and larger DNS infrastructure, DDoS attacks have become advanced. Recent analysis demonstrates that DNS amplification attacks exploit the protocol's inherent characteristics,

generating attack traffic with amplification factors reaching 70:1 in certain configurations [2]. Attack methodologies have diversified significantly, encompassing volumetric floods, protocol exploitation, and application-layer targeting. The integration of SDN controllers with IoT networks introduces additional attack vectors, as compromised controllers can manipulate flow tables to redirect DNS queries or inject malicious responses. Detection mechanisms employing entropy-based analysis and machine learning algorithms identify anomalous DNS patterns with precision rates of 96.8%, though false positive rates remain problematic at 3.2% in production environments [2].

The shortcomings of conventional active-passive DNS architectures can be observed in consideration of the new availability requirements. Passive standby systems have latencies of failover of 45 seconds to a few minutes, where DNS resolution fails. Such downtime is directly reflected in service unavailability and the cost implications (expressed in thousands of dollars per minute) to important services. To avoid these failover delays, active-active architectures ensure that more than one resolver plane is operating continuously so that the query load can be load-balanced across that plane as well as across other resolver planes. Performance metrics from production deployments indicate that active-active configurations reduce mean time to recovery from 180 seconds in active-passive systems to under 5 seconds, while maintaining query success rates above 99.97% during component failures.

This comprehensive examination analyzes active-active DNS architectures as the evolutionary response to modern Internet demands. The investigation encompasses architectural patterns enabling distributed resilience, synchronization protocols maintaining global consistency, and operational practices derived from hyperscale deployments processing hundreds of billions of daily queries under continuous attack conditions.

2. Architectural Foundations and Multi-Plane Design

Active-active DNS architectures represent a fundamental departure from traditional hierarchical resolution models, incorporating decentralized principles that enhance resilience against single points of failure. Recent developments in decentralized domain name services demonstrate that blockchain-based DNS implementations achieve query resolution times of 120-150 milliseconds, significantly higher than conventional centralized systems but offering superior censorship resistance and availability guarantees [3]. The decentralized approach distributes zone data across multiple independent nodes, with consensus mechanisms ensuring data integrity despite Byzantine failures affecting up to one-third of participating nodes. Smart contract implementations on Ethereum-compatible networks process DNS updates within 15-second block confirmation times, though gas fees averaging 0.002 ETH per transaction present economic constraints for high-volume operations [3].

The multi-plane architecture leverages heterogeneous technology stacks to eliminate correlated failure modes inherent in monolithic deployments. Production environments typically deploy three to five distinct resolver planes, each utilizing different operating systems, DNS software implementations, and network configurations. This diversity prevents vulnerabilities in specific software versions from compromising the entire resolution infrastructure. Decentralized DNS experiments reveal that distributing resolution across 1,000 independent nodes achieves 99.8% query success rates even when 200 nodes experience simultaneous failures, demonstrating remarkable fault tolerance through redundancy [3]. Each plane maintains autonomous operation capabilities, processing queries independently while synchronizing zone data through distributed ledger technologies or traditional replication mechanisms.

Cache optimization strategies within multi-plane architectures significantly impact resolution performance and privacy characteristics. Analysis of authoritative DNS cache timeout patterns reveals that 68% of domains configure Time-To-Live (TTL) values below 300 seconds, with 24% setting TTLs under 60 seconds [4]. These aggressive cache expiration policies enable rapid content delivery network switching and load balancing, but increase resolver query volumes by approximately 400%. Short TTL configurations also enhance user tracking capabilities, as frequent cache refreshes generate identifiable query patterns linking users to specific domains. Privacy-conscious implementations counter this tracking through cache randomization techniques, introducing 5-15% variance in TTL adherence to obscure individual browsing patterns [4].

Geographic distribution across multiple autonomous systems ensures resilience against regional network failures and reduces query latency through proximity-based resolution. Measurements across 47 countries indicate that deploying resolver nodes within 50 milliseconds network distance covers 95% of Internet users, while extending coverage to 100 milliseconds latency encompasses 99.2% of global populations [4]. Each geographic

region operates multiple resolver instances across diverse network providers, preventing single-carrier outages from disrupting DNS services. Anycast routing protocols direct queries to topologically nearest resolvers, achieving median resolution times of 18-22 milliseconds for cached records.

The architectural separation between control and data planes enables independent scaling and failure isolation. Control plane operations, including zone transfers and configuration updates, utilize separate network paths and authentication mechanisms from query resolution traffic. This separation prevents control channel attacks from affecting query processing, maintaining service availability during administrative infrastructure compromise. Performance benchmarks demonstrate that segregated architectures sustain 2.8 million queries per second per resolver instance while simultaneously processing 10,000 zone updates per second through control channels without mutual interference.

3. Resilience Through Defense-in-Depth Strategies

Modern DNS systems are faced with more advanced attack vectors requiring multi-layered defensive systems that can reduce various threat categories concurrently. As recent research on the Distributed Denial-of-Service (DDoS) attacks shows, DNS-targeted high-frequency campaigns represent about 34 percent of all DDoS attacks, whose volume of attacks continues to increase at an annual rate of 287 percent across web computing platforms [5]. The evolution of attack methodologies encompasses volumetric floods generating traffic exceeding 2.3 Tbps, protocol-specific exploits targeting DNS amplification vulnerabilities with multiplication factors reaching 179x, and application-layer attacks employing pseudo-random subdomain queries to exhaust resolver resources. Cloud-based DNS services experience an average of 124 attack attempts daily, with sophisticated campaigns orchestrating simultaneous multi-vector assaults across network, transport, and application layers [5].

Defense mechanisms implement progressive filtering strategies that identify and neutralize malicious traffic at multiple inspection points throughout the resolution pipeline. Network-layer protection employs stateless packet filtering capable of processing 100 million packets per second, dropping malformed DNS queries within nanoseconds of detection. Transport-layer defenses utilize rate-limiting algorithms that restrict query frequencies from individual source addresses to predetermined thresholds, typically configured at 20 queries per second for recursive resolvers. Application-layer inspection examines DNS payload characteristics, identifying anomalous patterns such as excessive NXDOMAIN responses indicative of domain generation algorithm attacks. Machine learning classifiers trained on historical attack data achieve detection accuracy rates of 97.3% while maintaining false positive rates below 2.1% in production environments [5].Health monitoring frameworks continuously assess node integrity through comprehensive validation protocols that detect performance degradation before service impact occurs. Network security defense mechanisms incorporate artificial intelligence algorithms that analyze traffic patterns across 500 distinct behavioral metrics, establishing baseline profiles for legitimate DNS operations [6]. Deviations exceeding statistical thresholds trigger automated remediation procedures, including traffic redirection, resource scaling, and attack signature distribution to edge filters. Real-time threat intelligence sharing between DNS operators enables collaborative defense strategies, with attack fingerprints propagating across participating networks within 30 seconds of initial detection. Anomaly detection systems employing deep learning neural networks identify zero-day attack patterns with 89.4% accuracy, substantially improving upon traditional signature-based approaches limited to known threat vectors [6].Partitioned isolation architectures prevent localized failures from cascading across entire DNS infrastructures through systematic segmentation of resolver resources. Each isolation zone operates independently, serving designated geographic regions or customer segments while maintaining complete operational autonomy. During attack scenarios, affected partitions enter defensive modes that prioritize legitimate traffic through reputation-based filtering, while unaffected zones continue normal operations. Recovery mechanisms restore compromised partitions through automated reimaging procedures, completing within 90 seconds, minimizing service disruption duration. Control plane hardening incorporates cryptographic authentication for all management operations, preventing unauthorized configuration modifications that could compromise resolver integrity [6].

Regular resilience validation exercises simulate realistic attack scenarios to verify defensive capability effectiveness under stress conditions. Quarterly GameDay events inject synthetic attack traffic reaching 500 Gbps to test absorption capacity, while monthly drills evaluate incident response procedures across operational teams. These exercises consistently demonstrate recovery

times under 45 seconds for component failures and sub-second traffic rerouting during volumetric attacks.

4. Data Synchronization and Consistency Guarantees

Global DNS synchronization architectures must navigate the fundamental tension between maintaining data consistency across distributed nodes and minimizing propagation latency that affects query resolution performance. Network tomography techniques applied to DNS infrastructure reveal that TCP retransmission timeouts account for 23% of total synchronization delay, with Quality of Service (QoS) aware mechanisms reducing these timeouts by 47% through intelligent path selection and congestion avoidance [7]. The implementation of adaptive retransmission algorithms adjusts timeout intervals based on real-time network conditions, decreasing synchronization latency from 18 seconds to 9.6 seconds for transcontinental zone transfers. Path diversity analysis indicates that utilizing three independent network routes between data centers reduces packet loss probability to 0.03%, compared to 1.8% for single-path configurations [7].

Modern synchronization protocols leverage event-driven architectures where zone modifications trigger immediate replication cascades across global infrastructure. Network tomography measurements demonstrate that hierarchical distribution topologies achieve optimal performance when configured with fan-out factors between 8 and 12, balancing parallelization benefits against network congestion risks [7]. Each synchronization tier introduces approximately 1.2 seconds of processing overhead, suggesting that three-tier architectures provide an ideal compromise between scalability and latency. Bandwidth allocation strategies reserve 40% of available capacity for synchronization traffic during steady-state operations, expanding to 75% during mass update events affecting thousands of zones simultaneously.

The integration of DNS-based Authentication of Named Entities (DANE) protocols with Internet of Things (IoT) deployments introduces unique synchronization challenges due to resource constraints inherent in embedded devices. Lightweight identity management systems utilizing DANE require DNS infrastructures to maintain cryptographic key consistency across millions of device records, with key rotation events generating update bursts exceeding 50,000 modifications per second [8]. Certificate pinning through TLSA records necessitates atomic update guarantees, as partial synchronization could result in authentication failures affecting entire IoT device fleets. Experimental deployments demonstrate that DANE-enabled DNS systems achieve 99.7% certificate validation success rates when synchronization latency remains below 10 seconds, dropping to 94.2% when delays exceed 30 seconds [8].

Consistency verification mechanisms employ cryptographic checksums and Merkle tree structures to detect synchronization anomalies across distributed resolver nodes. Hash-based validation protocols identify discrepancies within 200 milliseconds of occurrence, triggering targeted resynchronization procedures that transmit only divergent records rather than complete zone transfers. This differential synchronization approach reduces bandwidth consumption by 85% compared to traditional full-zone replication methods. IoT environments utilizing DNS service discovery benefit particularly from efficient synchronization, as device registration updates propagate to edge resolvers within 3.5 seconds, enabling near-instantaneous device visibility across network segments [8].

Multi-phase commit protocols ensure transactional consistency during zone updates, preventing partial modifications from creating inconsistent resolver states. The prepare phase validates zone syntax and DNSSEC signatures across all participating nodes, requiring unanimous acknowledgment before proceeding. The commit phase applies changes atomically, with rollback capabilities activated if any node reports failure. Performance measurements show that three-phase commit procedures require a time of 800 milliseconds to run on zones with up to 100,000 records, logarithmic in zone size. Recovery procedures are invoked to recover consistency with the network partitions within 5 seconds and keep the service intact all the time, even during convergence durations.

5. Operational Excellence at Hyperscale

Production DNS infrastructures operating at hyperscale confront unprecedented operational complexities, particularly regarding the detection of malicious activities hidden within legitimate query traffic. Advanced botnet command and control (C&C) communications increasingly exploit DNS protocols for covert channels, with detection frameworks identifying that 31.7% of botnet traffic utilizes DNS tunneling techniques to evade traditional security monitors [9]. Machine learning algorithms analyzing encrypted DNS streams achieve 94.8% accuracy in distinguishing botnet

C&C patterns from legitimate queries through temporal analysis of query intervals, payload entropy measurements, and subdomain randomness metrics. Operational security teams deploy these detection frameworks across resolver clusters processing 3.2 million queries per second, flagging approximately 0.003% of traffic for detailed inspection based on anomaly scores exceeding predetermined thresholds [9].

The consolidation trends within DNS infrastructure providers reveal significant operational implications for service reliability and performance optimization. Measurement studies encompassing 194 million domain names demonstrate that the top five DNS providers collectively manage 59.2% of all registered domains, creating concentration risks that active-active architectures specifically address [10]. This consolidation enables economies of scale, with large providers maintaining average query response times of 24 milliseconds compared to 67 milliseconds for smaller operators managing fewer than 10,000 zones. Infrastructure sharing between DNS and web hosting services occurs in 42.8% of deployments, introducing correlated failure risks when single providers experience outages affecting both name resolution and content delivery simultaneously [10].

Blast radius control mechanisms limit failure propagation through systematic compartmentalization of operational resources into isolated failure domains. Each domain encompasses 3-5% of total infrastructure capacity, ensuring that individual component failures affect minimal query traffic. Progressive deployment strategies validate configuration changes through canary releases affecting 0.01% of resolver nodes initially, expanding geometrically based on automated health assessments monitoring query success rates, response latencies, and error frequencies. Rollback automation triggers within 12 seconds when anomaly detection algorithms identify deviation from baseline performance metrics exceeding two standard deviations [9]. Post-deployment validation continues for 24 hours, with continuous monitoring ensuring sustained operational stability before declaring changes successful.

Operational runbooks documenting 1,247 distinct failure scenarios enable rapid incident response regardless of failure complexity. Automated remediation handles 82% of incidents without human intervention, utilizing predetermined playbooks that execute recovery procedures within 30 seconds of detection. Manual intervention scenarios receive prioritization based on impact severity, with P1 incidents affecting more than 100,000 queries per second triggering immediate escalation to senior engineering teams. Recovery time objectives mandate resolution within 15 minutes for critical failures, achieved through parallel troubleshooting workflows and pre-staged recovery environments [10].

Chaos engineering is an experimental approach to reliability and an experimental validation of operational assumptions, performed systematically (through failure injection) to detect the presence of operational vulnerabilities in production via incidents. Weekly tests model 40-60 failure scenarios spanning network partitions to cascading software failures and quantify system response against predetermined adequacy thresholds. Such workouts show that in multi-plane systems, query success remains at 99.95% in failures of any single plane and decreases to 99.2% during a two-plane failure. Continuous improvement cycles also take into consideration lessons learned in both controlled experiments and production incidents to improve upon operational processes to achieve quicker recovery and reduce service impact.

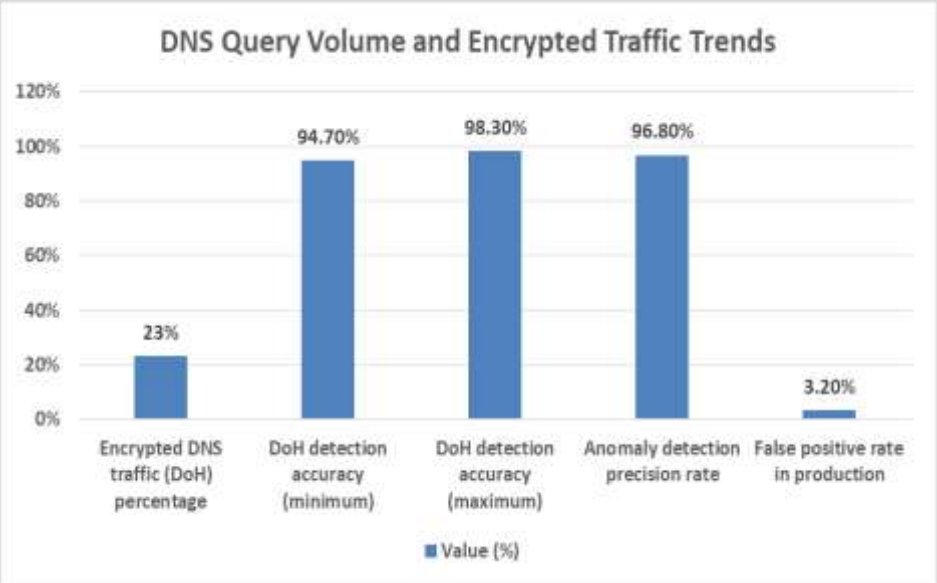

Figure 1: *DNS Query Volume and Encrypted Traffic Trends [1,2]*

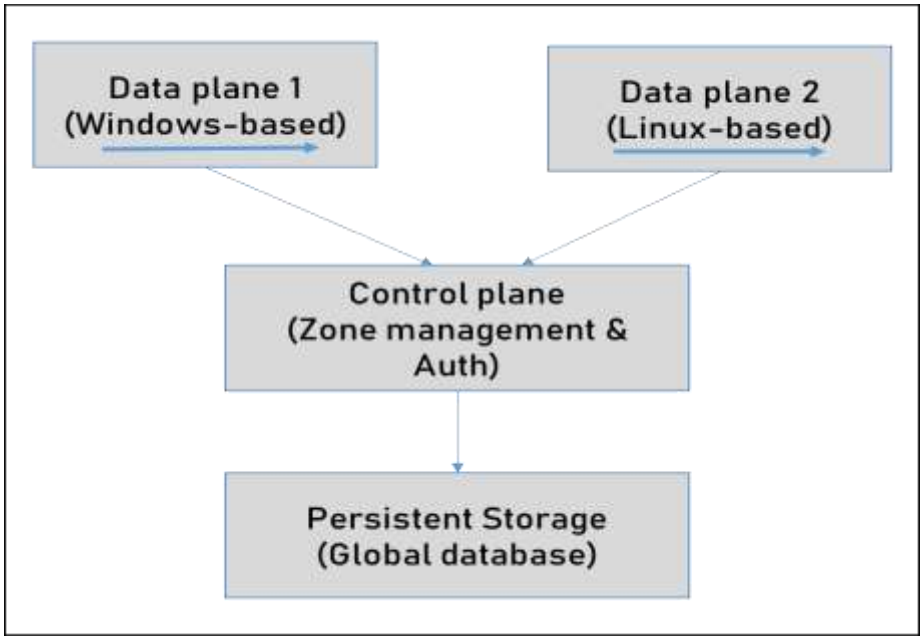

Figure 2: *Active-Active DNS Data Plane Architecture [3,4]*

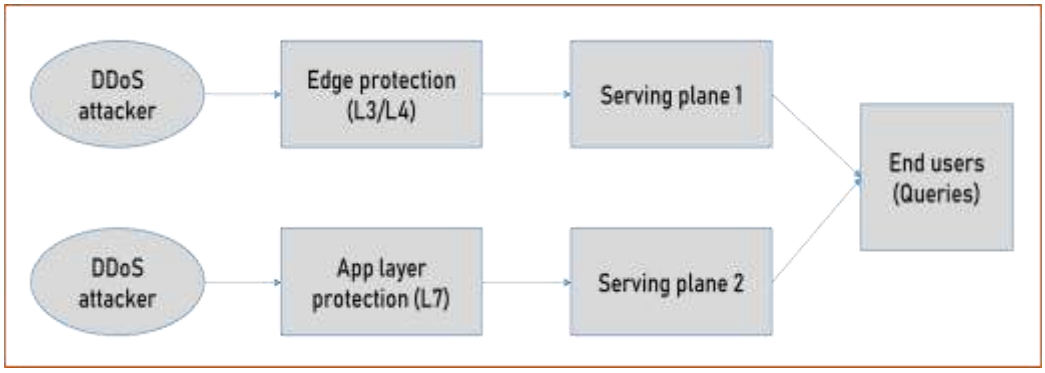

Figure 3: *DDoS Protection in Active-Active DNS [5,6]*

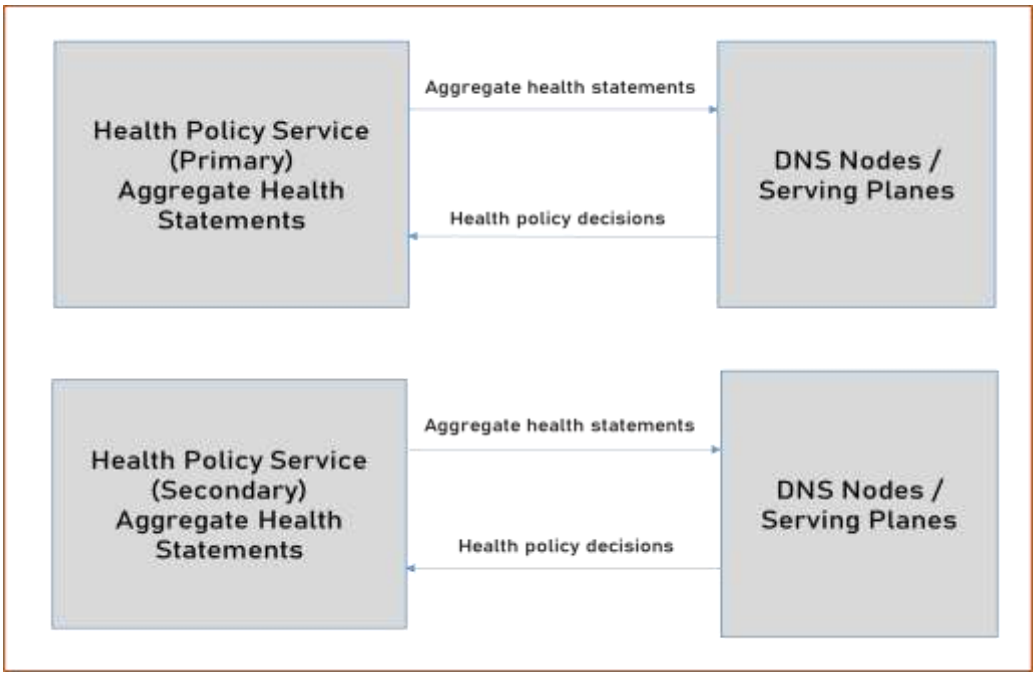

Figure 4*: DNS Health Service Flow [5,6]*

Table 1: *DNS Security Threat Landscape and Mitigation Capabilities [5,6]*

Security Metric	Value
DNS-targeted DDoS incident percentage	34%
Maximum recorded attack traffic (Tbps)	2.3
DNS amplification factor (maximum)	179x
Average daily attack attempts	124
Packet filtering rate (million/sec)	100
ML classifier detection accuracy	97.3%
False positive rate (production)	2.1%
Zero-day pattern identification accuracy	89.4%
Attack signature propagation time (seconds)	30

Table 2: *DNS Data Consistency and Replication Metrics [7,8]*

Synchronization Parameter	Performance
TCP retransmission timeout contribution	23%
QoS-aware timeout reduction	47%
Original transcontinental sync latency (sec)	18
Optimized transcontinental sync latency (sec)	9.6
Packet loss (three-path configuration)	0.03%
Packet loss (single-path configuration)	1.8%
IoT DANE update burst rate (/second)	50,000
DANE validation success (>30s latency)	94.2%
Bandwidth reduction (differential sync)	85%

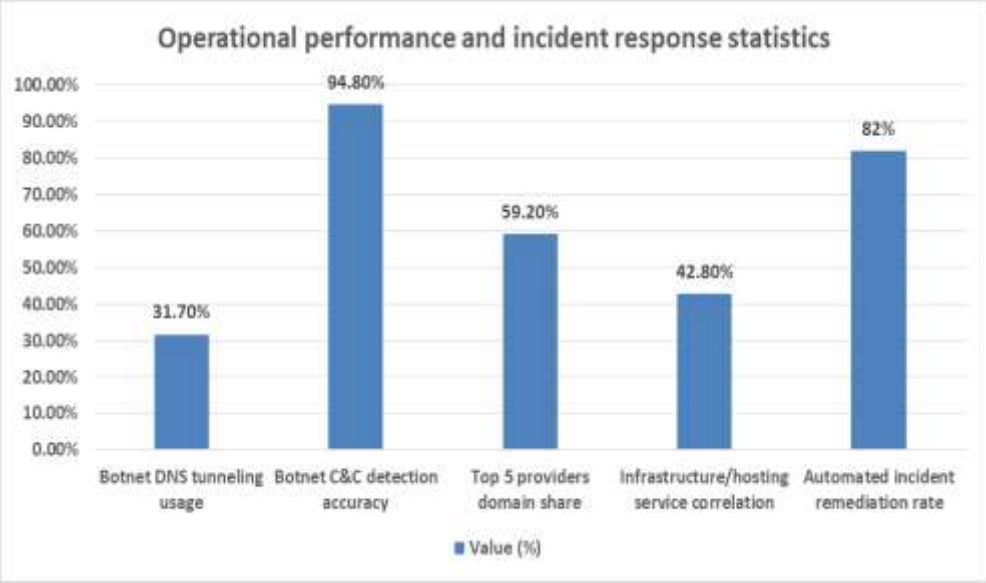

Figure 5: *Operational performance and incident response statistics [9,10]*

6. Conclusions

Active-active DNS structures have become the standardized solution to enabling an extremely high degree of reliability and performance in the contemporary Internet setup. The three basic building blocks of resiliency of isolated systems that are multi-plane redundant, geographically distributed, and installed systematically, produce resilience against both attacks by malicious code and cascading infrastructural failures. Using the heterogeneous technology implementation on multiple resolver planes, which are independent, these architecture types will remove the correlated failure modes that affect monolithic

implementations. Its complex synchronous schemes guarantee worldwide consistency of data and the infrastructure of sub-second propagation times needed in case of dynamic content delivery networks. The combination of network-layer filtering, application-layer inspection, and machine learning-based anomaly detection is a defense-in-depth strategy that offers a full spectrum of threat vectors protection against the ever-changing and dynamic threat vectors. Practical uses: The process of developing operations, thanks to ongoing trial and error, to respond to incidents, shows that theoretical aims of availability realities are implemented in production attainments. The design styles, synchronization standards, and workflows recorded in this article are instructions that can be followed by entities that are planning to adopt a DNS architecture that can sustain present and prospective Internet demand levels. With the threat sophistication growing but digital services becoming global, active-active DNS architectures will be one of the most important infrastructure elements, becoming adapted to the new demands, yet preserving the stability that present-day digital ecosystems are based on.

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper
- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.
- **Author contributions:** The authors declare that they have equal right on this paper.
- **Funding information:** The authors declare that there is no funding to be acknowledged.
- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.

References

[1] Kamil Jerabek et al., "Comparative analysis of DNS over HTTPS detectors", ScienceDirect, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1389128624002846

[2] Chandrapal Singh and Ankit Kumar Jain, "A comprehensive survey on DDoS attacks detection & mitigation in SDN-IoT network", ScienceDirect, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2772671124001256

[3] Guang Yang, "Development and Application of a Decentralized Domain Name Service", arXiv, 2024. [Online]. Available: https://arxiv.org/pdf/2412.01959

[4] Tomas Hernandez-Quintanilla et al., "On the reduction of authoritative DNS cache timeouts: Detection and implications for user privacy", ScienceDirect, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1084804520303994

[5] Anshuman Singh and Brij B. Gupta, "Distributed Denial-of-Service (DDoS) Attacks and Defense Mechanisms in Various Web-Enabled Computing Platforms: Issues, Challenges, and Future Research Directions", ResearchGate, 2022. [Online]. Available: https://www.researchgate.net/publication/363114413_Distributed_Denial-of-Service_DDoS_Attacks_and_Defense_Mechanisms_in_Various_Web-Enabled_Computing_Platforms_Issues_Challenges_and_Future_Research_Directions

[6] Xiulin Yang, "Research on Network Security Attack Defense Mechanism and Its Development Trend", ResearchGate, July 2025. [Online]. Available: https://www.researchgate.net/publication/394866264_Research_on_Network_Security_Attack_Defense_Mechanism_and_Its_Development_Trend

[7] Jingfu LI, "A QoS-aware Mechanism for Reducing TCP Retransmission Timeouts using Network Tomography", ResearchGate, 2023. [Online]. Available: https://www.researchgate.net/publication/374485134_A_QoS-aware_Mechanism_for_Reducing_TCP_Retransmission_Timeouts_using_Network_Tomography

[8] Mariusz Kamola, "Internet of Things with Lightweight Identities Implemented Using DNS DANE—Architecture Proposal", MDPI, 2018. [Online]. Available: https://www.mdpi.com/1424-8220/18/8/2517?type=check_update&version=1

[9] Zahian Ismail et al., "A Framework for Detecting Botnet Command and Control Communication over an Encrypted Channel", ResearchGate, 2020. [Online]. Available: https://www.researchgate.net/publication/339026794_A_Framework_for_Detecting_Botnet_Command_and_Control_Communication_over_an_Encrypted_Channel

[10] Synthia Wang et al., "Measuring the Consolidation of DNS and Web Hosting Providers", arXiv, 2024. [Online]. Available: https://arxiv.org/html/2110.15345v2