Journal of Information Systems Engineering and Management
2026, 11(2s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ **Research Article**

Active-Active DNS Resiliency Across Two Providers: Architecture, Performance, and Operational Efficiency

Surya Narayana Lankalapalli

Microsoft Corporation, USA

ARTICLE INFO

Received: 20 Dec 2025

Revised: 06 Feb 2026

Accepted: 16 Feb 2026

ABSTRACT

Digital service delivery has undergone substantial transformation as organizations extend their operational footprint across international markets. Contemporary users anticipate seamless connectivity independent of geographic position, network volatility, or underlying infrastructure complications. Conventional single-provider or Active-Passive DNS failover methodologies demonstrate increasing inadequacy when addressing these requirements, primarily attributed to record propagation latencies, initialization challenges, and concentrated provider reliance. Active DNS implementation utilizing dual autonomous DNS providers resolves these constraints through concurrent traffic servicing from multiple operational endpoints while preserving provider-tier redundancy. Rather than establishing a primary DNS provider designation with secondary backup infrastructure, both providers simultaneously respond to queries, allocate traffic loads, and execute routing determinations. This dual-provider Active-Active framework achieves near-instantaneous recovery intervals, eliminates DNS provider dependency, and advances international performance through proximity-driven routing. Additionally, this article guarantees that complete deployed infrastructure participates in production operations rather than maintaining dormant status. This article investigates how dual-provider Active-Active DNS reinforces availability parameters, advances performance characteristics, and strengthens operational effectiveness through sophisticated routing mechanisms and resilient architectural frameworks.

Keywords: Active-Active DNS, DNS resilience, multi-provider architecture, traffic management, global availability

1. Introduction

1.1 Evolution of Requirements and Service Expectations Globally

The Domain Name System constitutes fundamental infrastructure enabling internet functionality, performing translation between human-interpretable domain identifiers and numerical IP addresses that facilitate service accessibility [1]. Organizations have progressively expanded their digital infrastructure across numerous geographic territories, fundamentally altering availability expectations from scheduled maintenance intervals toward perpetual operational states. Contemporary users engage with digital platforms across heterogeneous network environments, spanning high-capacity fiber connections to mobile networks exhibiting variable latency characteristics and reliability profiles. This transformation has substantially modified infrastructure design philosophies regarding DNS implementation, transitioning from concentrated single-region deployments toward internationally distributed architectures capable of sustaining service continuity independent of localized failures or network segmentation events. The pivotal position DNS occupies within service delivery architectures indicates that transient DNS disruptions can precipitate substantial revenue deterioration and customer satisfaction degradation.

1.2 Constraints Within Conventional DNS Frameworks

Conventional single-provider DNS configurations establish intrinsic vulnerability patterns despite operational straightforwardness. When organizations rely on only one DNS provider, they also accept the risks of that provider's infrastructure failure modes, capacity limits, and network routing decisions. Active-Passive failover frameworks attempt to address these issues through standby DNS infrastructure maintenance that activates during primary system disruptions. Nevertheless, these methodologies introduce distinct challenges, including TTL-driven propagation delays potentially preventing users from resolving domain identifiers for extended durations, initialization complications where backup systems must rapidly scale under complete production loading, and economic inefficiency associated with maintaining fully provisioned infrastructure remaining dormant during standard operations. Service Level Agreements have become important tools for measuring these limits. They set measurable standards for performance, availability, and reliability that businesses use to judge how their DNS infrastructure is [2].

1.3 Dual-provider Active-Active Framework

Active-Active DNS across dual autonomous DNS providers represents fundamental resilience strategy transformation. Rather than designating primary and backup system roles, this architectural approach treats both providers as equivalent participants in traffic administration. Each provider maintains proprietary global network infrastructure, executes independent health surveillance, and renders autonomous routing determinations based on real-time endpoint status information. When appropriately configured, both providers continuously service production traffic, eliminating idle failover capacity concepts. Recent investigations have demonstrated substantial improvements in DNS optimization beyond conventional content delivery networks, indicating that active routing strategies can meaningfully enhance service delivery across heterogeneous geographic territories and network conditions [3]. This methodology transforms DNS from a potential singular failure point into a distributed architecture where individual component outages cannot interrupt global service availability.

1.4 Examination Scope and Structural Overview

This examination evaluates technical components, performance characteristics, and operational advantages associated with dual-provider Active-Active DNS architectures. The analysis encompasses health check mechanisms enabling both providers to render independent routing determinations, intelligent routing policies optimizing traffic distribution based on latency and geography, TTL management strategies balancing responsiveness with stability, performance optimization techniques reducing user-perceived latency, and the cost implications of comprehensively utilizing deployed infrastructure. The implementation of Active-Active architectures in global-scale name resolution systems has demonstrated that resilient DNS configurations can maintain service continuity even during substantial infrastructure disruptions, providing a validated foundation for organizations requiring maximum availability [4].

2. Architectural Foundations and Resiliency Principles

2.1 Comparative Evaluation of Resiliency Methodologies

DNS resiliency strategies have progressed through multiple generations, each addressing preceding approach limitations. Single-provider configurations offer operational simplicity but establish provider dependency and singular failure points. Active-Passive failover introduces redundancy but experiences detection delays, TTL propagation lag, and risks that failover systems may not perform identically to primary infrastructure under production loading. Multi-provider Active-Active architectures eliminate these compromises by continuously distributing traffic across both providers, ensuring routing logic remains active and validated continuously. When an endpoint becomes unhealthy, both providers independently detect the failure and remove the affected target from DNS responses without requiring coordination or failover procedures. This architectural methodology aligns with broader industry trends toward distributed systems that eliminate singular failure points rather than attempting to construct perfectly reliable individual components.

Characteristic	Single-Provider	Active-Passive	Multi-Provider Active-Active
provider Dependency	High	Medium	Low
Infrastructure Utilization	Moderate	Low (40-60%)	High (80-95%)
Recovery Time Objective	Minutes to Hours	Minutes	Seconds
Failover Mechanism	Manual or Automatic	Automatic Detection Required	Continuous Self-Healing
Cold-Start Issues	Not Applicable	Present	Eliminated
TTL Propagation Impact	Moderate	High	Minimal
Configuration Complexity	Low	Medium	High
Operational Testing	Limited	Periodic	Continuous
Cost Efficiency	Moderate	Low	High

Table 1: Comparison of DNS Resiliency Models [1, 2, 3]

2.2 Provider Independence and Isolation of Faults

The dual-provider framework establishes meaningful fault isolation at the DNS infrastructure stratum itself. Each provider operates independent networks of authoritative name servers distributed internationally, utilizes distinct routing algorithms and network peering relationships, maintains separate monitoring and health verification systems, and responds to DDoS attacks or infrastructure failures without affecting the alternative provider. This independence ensures that complications affecting one DNS network, whether attributable to software defects, configuration errors, capacity exhaustion, or malicious attacks, do not propagate to the alternative provider. Users whose recursive resolvers are configured to query both authoritative name server sets can seamlessly fail over between providers at the query level, often without perceiving any service disruption. Advanced security mechanisms in contemporary DNS frameworks have incorporated machine learning-based threat detection systems that operate in real-time, enhancing the reliability of health inspection processes and improving the overall security posture of Active-Active deployments [5].

2.3 Recovery Intervals and Availability Parameters

Traditional availability metrics focus on system uptime percentages, but Active-Active DNS architectures shift emphasis toward Recovery Time Objectives, measuring how rapidly traffic reroutes around failures. In Active-Passive systems, RTO includes failure detection time, DNS record update propagation, TTL expiration at recursive resolvers, and potential initialization delays at backup infrastructure. These factors can accumulate to create recovery windows measured in minutes or extended durations. Active-Active systems compress RTO toward the health examination interval plus one TTL period, eliminating failover coordination time since both providers remain perpetually active, removing initialization delays since all infrastructure continuously processes production traffic, and avoiding DNS propagation delays since both providers independently update records. Organizations implementing these architectures often achieve effective RTOs measured in seconds rather than minutes, fundamentally modifying the user experience during infrastructure failures.

2.4 Distribution Patterns and Load Distribution

With two DNS providers operating in Active-Active mode, traffic distribution becomes a function of how recursive resolvers select among multiple authoritative name servers. Most resolvers use heuristics that take into account response time, availability, and random selection to decide which authoritative name server to ask. Over time, this procedure creates roughly balanced traffic distribution across both providers, though short-term variations occur due to network conditions and resolver behavior. Both providers independently implement routing policies that direct users to healthy endpoints based on geographic proximity, current latency measurements, or configured weights. Geolocation routing capability has been extensively studied and implemented in authoritative DNS servers, demonstrating practical benefits for directing users to regionally appropriate endpoints and reducing overall latency in internationally distributed applications [6]. This dual-layer optimization means that even if one provider's routing algorithm makes suboptimal decisions for a particular user population, the alternative provider's independent routing logic provides a second opportunity for optimal traffic placement.

3. Technical Elements and Deployment Considerations

3.1 Health Verification and Endpoint Surveillance

Health checks form the foundation of Active-Active DNS reliability. Each provider must independently verify that endpoints can successfully process user requests before including them in DNS responses. Effective health check implementations combine multiple verification layers, including TCP connection establishment to confirm network reachability, HTTP or HTTPS requests to validate application functionality, and application-specific probes that verify backend systems and dependencies. Both providers must implement equivalent health assessment logic to ensure consistent routing decisions, though complete synchronization is neither required nor desirable since independent verification provides additional fault isolation. Health evaluation intervals must balance responsiveness against backend load, with typical configurations polling endpoints every few seconds to detect failures quickly while avoiding overwhelming monitored systems with health examination traffic.

Parameter	Typical Range	Recommended Value	Impact on System
Health Check Interval	1-30 seconds	3-5 seconds	Faster detection vs. increased load
Failure Threshold	1-5 consecutive failures	2-3 failures	Balance between sensitivity and stability
Timeout Duration	1-10 seconds	2-3 seconds	Quick failure detection vs. false positives
Protocol Type	TCP, HTTP, HTTPS, Custom	HTTP/HTTPS with path validation	Application-level verification accuracy
Success Criteria	Status code, response time	200 status + sub-second response	Comprehensive endpoint health validation
Recovery Validation	1-5 consecutive successes	2 successes	Prevents flapping and instability
Regional Distribution	Single, Multiple regions	Multiple global vantage points	Geographic failure isolation capability

Table 2: Health Check Configuration Parameters [4, 5]

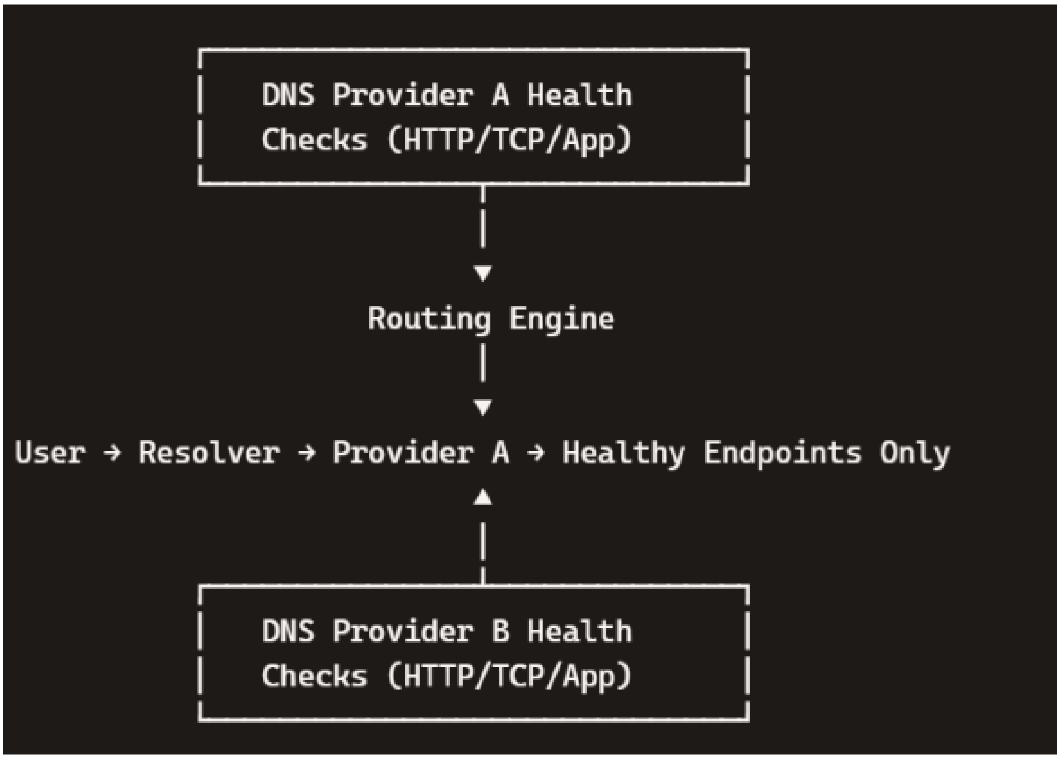

Fig. 1: Active-Active DNS Health Check Architecture and Traffic Flow

3.2 Sophisticated Routing Mechanisms Across providers

Advanced routing policies enable both DNS providers to direct users to optimal endpoints based on multiple factors. Latency-based routing checks the round-trip time of the network from different points of view to each endpoint and sends queries to the fastest available option. It automatically adjusts to changes in network conditions and traffic patterns. Geographic routing uses the source IP address of recursive resolvers to estimate user location and returns geographically proximate endpoints, reducing network distance and improving performance for internationally distributed applications. Weighted routing assigns traffic percentages to different endpoints, enabling gradual traffic shifts during deployments, capacity adjustments, or controlled experiments. Running these policies across two independent DNS networks provides redundancy in routing logic itself, ensuring that no single provider's routing algorithm becomes a singular failure point for traffic distribution decisions.

3.3 Time-to-Live Administration and Propagation Control

Time-to-Live values control how long recursive resolvers cache DNS responses before querying authoritative name servers again. In Active-Active architectures, TTL selection significantly impacts failure detection time and system responsiveness. Short TTLs enable rapid reaction to failures since cached records expire quickly, allowing resolvers to obtain updated DNS responses that reflect current endpoint health. However, short TTLs increase query volume against authoritative name servers and reduce caching benefits across the global resolver infrastructure. In-network caching strategies on programmable switches have demonstrated that DNS service optimization can significantly improve query response times and overall system efficiency, suggesting that intelligent caching at multiple network layers can partially offset the cost of shorter TTLs [7]. Typical Active-Active deployments use TTL values between thirty and one hundred twenty seconds to balance responsiveness with query load. Both providers must implement consistent TTL values to ensure predictable behavior and avoid situations where different resolvers receive responses with dramatically different caching durations.

3.4 Configuration Alignment and Deployment Automation

While Active-Active DNS providers operate independently, their configurations must remain aligned to ensure consistent routing behavior. Most organizations have configuration management systems that set the health check parameters for endpoints from both providers, the routing policy definitions and traffic weights for each endpoint, the TTL values for all DNS record types, and the geographic mappings for location-based routing. Automated deployment pipelines propagate configuration changes to both providers simultaneously or in rapid succession, reducing the window where providers might have divergent configurations. However, perfect synchronization is neither achievable nor necessary since independent decision-making provides valuable fault isolation. The objective is operational consistency rather than lock-step coordination, allowing each provider to independently respond to failures while maintaining similar routing logic under normal conditions.

4. Performance Enhancement and User Experience Optimization

4.1 Latency Minimization Through Distributed Networks

Active-Active DNS across two providers improves performance by providing users multiple optimized paths to reach the nearest healthy endpoint. Each provider operates proprietary global networks of authoritative name servers distributed across dozens of points of presence worldwide. When a recursive resolver asks for DNS records, it usually picks the authoritative name server that responds the fastest. This is usually the server that is closest in terms of physical or topological distance. By utilizing two providers, organizations effectively double the number of points of presence where their DNS records are served, increasing the probability that any given user has a nearby authoritative name server available. Performance benchmarking of authoritative DNS servers has provided methodologies for measuring and comparing DNS server capabilities under various load conditions, establishing baselines for evaluating Active-Active configurations and quantifying the latency benefits of multi-provider deployments [8]. This geographic distribution reduces DNS resolution time, which directly impacts page load times, API response latency, and overall user experience.

Routing Policy	Primary Benefit	Latency Reduction	Complexity	Best Application
Latency-Based	Fastest response time	30-50%	High	Real-time collaboration tools
Geographic	Regional proximity	20-40%	Medium	Content localization services
Weighted	Traffic control	Variable	Low	Gradual deployments and testing
Failover	Availability	Not applicable	Low	Backup endpoint designation
Geoproximity	Distance optimization	25-45%	Medium	Global application distribution
Multi-Value	Load distribution	10-20%	Low	Simple round-robin needs
Hybrid (Combined)	Comprehensive optimization	40-60%	Very High	Enterprise-scale applications

Table 3: Routing Policy Performance Characteristics [3, 6, 8]

4.2 Tolerance to Network Congestion and Path Complications

Internet routing demonstrates inherent unpredictability, with path changes, congestion events, and peering disputes regularly affecting network performance. When DNS resolution relies on a single provider's network, users experience whatever performance that network delivers independent of whether superior paths exist. Active-Active deployments provide automatic path diversity since recursive resolvers can independently query either provider's authoritative name servers. If one provider's routes to a certain area are congested, resolvers in that area will naturally switch to the other provider's infrastructure, which probably uses different network paths and peering relationships. This automatic adaptation occurs without manual intervention or configuration changes, as resolvers simply favor the authoritative name servers that provide faster responses. The result is more consistent DNS resolution performance across diverse network conditions and geographic regions.

4.3 Workload Performance for Latency-Critical Applications

Certain application categories experience disproportionate benefits from optimized DNS resolution. Financial trading platforms require minimal latency since even small delays can impact transaction execution and market opportunity. Real-time collaboration tools depend on quick connection establishment to maintain smooth user interactions and avoid perceptible lag. Interactive APIs serve mobile and web applications where users expect immediate responses, making every millisecond of latency visible in the user experience. For these workloads, Active-Active DNS provides meaningful performance improvements by ensuring DNS resolution occurs as quickly as possible from the nearest available provider, maintaining consistent performance even during regional network complications, and eliminating DNS-related delays during provider infrastructure problems. Organizations operating latency-sensitive services often consider DNS optimization a critical component of overall performance strategy rather than an afterthought.

4.4 Deployment Risk Reduction and Traffic Control

Using weighted routing with two DNS providers helps make application launches safer by letting organizations slowly move traffic between different versions or environments. Organizations can begin directing a small percentage of traffic to new infrastructure while monitoring error rates, performance metrics, and user feedback. If complications arise, traffic weights can be instantly adjusted to redirect users back to known-good infrastructure without waiting for new deployments or configuration rollbacks. Both providers independently implement these weights, providing redundancy in traffic control mechanisms. This approach reduces deployment risk by validating changes under real production traffic before full rollout, enabling rapid rollback if problems emerge, and maintaining fine-grained control over traffic distribution throughout deployment processes. The operational confidence gained from safe deployment mechanisms often proves as valuable as the technical resilience benefits of Active-Active architectures.

5. Operational Effectiveness and Financial Considerations

5.1 Infrastructure Utilization and Cost Efficiency

Active- DNS ensures all deployed infrastructure actively serves production traffic rather than sitting idle as standby capacity. In Active-Passive models, organizations must provision full backup capacity that provides no value during normal operations—a significant inefficiency for globally distributed deployments.

Active-Active architectures eliminate this waste: every provisioned endpoint processes real traffic continuously, and capacity planning reflects actual utilization rather than theoretical peak loads.

Organizations typically achieve substantially higher infrastructure utilization compared to Active-Passive deployments, translating directly to cost savings [9].

Cost Component	Active-Passive	Active-Active	Savings
Infrastructure Provisioning	200% (Primary + Standby)	100% (All Active)	50%
Emergency Scaling Events	Frequent	Rare	30-40%
Downtime Revenue Impact	High	Minimal	70-90%
Total Cost of Ownership	Baseline	75-85% of Baseline	15-25%

Table 4: Infrastructure Cost Comparison [9]

5.2 Operational Complexity Trade-offs

Active-Active architectures introduce additional complexity—maintaining dual provider relationships, ensuring configuration consistency, and monitoring both systems. However, this is offset by reduced incident response pressure (failures self-heal automatically), elimination of manual failover procedures, and better visibility through cross-provider telemetry comparison. Modern infrastructure-as-code practices significantly reduce the manual overhead of multi-provider operations.

5.3 Security Considerations

The distributed nature of dual-provider architectures provides inherent DDoS resilience since attackers must overwhelm both providers simultaneously. While the expanded attack surface requires careful security planning [10], the overall security posture generally improves: compromising DNS resolution requires breaching both providers, and provider network independence limits the blast radius of successful attacks.

6. Business Value and Service Continuity Implications

6.1 Revenue Protection and Customer Trust

DNS outages directly impact revenue through interrupted transactions, abandoned sessions, and failed API integrations. More critically, frequent outages erode customer confidence and drive users toward competitors. Active-Active architectures protect both revenue streams and brand reputation by maintaining availability during provider failures and eliminating the service gaps associated with failover delays.

6.2 Global Expansion and Compliance

For organizations expanding internationally, Active-Active DNS ensures new regions benefit from optimal routing immediately, with consistent performance across all markets regardless of regional infrastructure maturity.

The architecture also supports compliance requirements in regulated industries (financial services, healthcare, government) by enabling documented RTO/RPO metrics, demonstrating continuously-tested failover mechanisms, and providing audit trails of routing decisions.

Conclusion

Active-Active DNS across two independent providers represents a contemporary, resilient approach to global traffic management. By serving traffic simultaneously across multiple healthy endpoints and multiple DNS networks, organizations eliminate failover delays, improve performance, and fully utilize their infrastructure. The architecture addresses fundamental limitations of traditional DNS designs by removing singular failure points at the provider level, compressing recovery time toward zero through continuous health monitoring and automatic traffic rerouting, and enabling sophisticated routing policies that optimize user experience across diverse geographic regions and network conditions.

When combined with robust health checks, intelligent routing policies, and disciplined TTL management, multi-provider Active-Active DNS delivers near-continuous availability and operational flexibility at a global scale. Rather than treating resilience as a reactive process involving failure detection and recovery procedures, this architecture embeds fault tolerance directly into everyday traffic flow. Both DNS providers continuously participate in traffic management, ensuring that routing logic remains active and validated constantly rather than waiting dormant until needed.

The benefits extend beyond technical resilience to encompass economic efficiency through full infrastructure utilization, operational confidence through safe deployment mechanisms, and business value through consistent availability that protects revenue and customer trust. Organizations implementing Active-Active architectures must accept additional complexity in managing dual provider relationships and maintaining configuration consistency, but these costs are typically far outweighed by the improvements in availability, performance, and operational outcomes.

In an always-on digital environment where users interact with services across diverse devices, networks, and geographic locations, the objective is no longer simply recovering from failure quickly. Rather, contemporary architectures aim to ensure users never perceive failures at all, maintaining seamless service continuity independent of underlying infrastructure complications. Multi-provider Active-Active DNS makes that level of resilience achievable by transforming DNS from a potential singular failure point into a distributed system where no individual component outage can interrupt global service availability. As digital services continue to expand in reach and importance, architectural patterns that eliminate rather than merely mitigate singular failure points will become increasingly essential for organizations committed to meeting user expectations for continuous, high-performance access.

References

[1] Michael Dooley and Timothy Rooney, "Introduction to the Domain Name System (DNS)," IEEE Xplore, 2017. Available: https://ieeexplore.ieee.org/document/8008719?utm_source=copilot.com

[2] DN.org Staff, "Assessing DNS Service Level Agreements: Quantifying Performance, Availability, and Reliability," DN.org, April 17, 2025. Available: https://dn.org/assessing-dns-service-level-agreements-quantifying-performance-availability-and-reliability-in-a-critical-internet-layer/?utm_source=copilot.com

[3] Yue Wang et al., "ActiveDNS: Is There Room for DNS Optimization Beyond CDNs?," IEEE Xplore, 09 September 2024. Available: https://ieeexplore.ieee.org/document/10639696?utm_source=copilot.com

[4] Anil Puvvadi, "Active-Active DNS Architectures: Building Resilient Global-Scale Name Resolution Systems," International Journal of Computational and Experimental Science and Engineering (IJCESEN), October 11, 2025. Available: https://mail.ijcesen.com/index.php/ijcesen/article/download/4099/1283/10203

[5] Sanket Kolte et al., "A Machine Learning-Based Framework for Real-Time DNS Threat Detection and Mitigation Using Ensemble Models and Advanced Security Mechanisms," IEEE Xplore, 11 August

2025. Available: https://ieeexplore.ieee.org/document/11101638?utm_source=copilot.com

[6] Ta-Li Lai and Meng-Hsun Tsai, "Design and Implementation of a DNS Server with Geolocation Routing Capability," IEEE Xplore, 14 October 2021. Available: https://ieeexplore.ieee.org/document/9562605?utm_source=copilot.com

[7] Fan Yang et al., "DNS Service Optimization Through In-Network Caching on Programmable Switches," IEEE Xplore, 2025. Available: https://ieeexplore.ieee.org/document/10941754?utm_source=copilot.com

[8] Gábor Lencse et al., "Benchmarking Authoritative DNS Servers," IEEE Xplore, 14 July 2020. Available: https://ieeexplore.ieee.org/document/9139929?utm_source=copilot.com

[9] DN.org Staff, "DNS and Cloud Infrastructure Cost Optimization in Enterprise Environments," DN.org, March 31, 2025. Available: https://dn.org/dns-and-cloud-infrastructure-cost-optimization-in-enterprise-environments/?utm_source=copilot.com

[10] Anju Ramdas et al., "A Survey on DNS Security Issues and Mitigation Techniques," IEEE Xplore, 16 April 2020. Available: https://ieeexplore.ieee.org/document/9065354?utm_source=copilot.com