Alqalam Journal of Medical and Applied Sciences. 2026;9(6):1522-1526
https://doi.org/10.54361/ajmas.269610

Systematic Review

IPv4 Address Exhaustion: Mitigation Strategies, Architectural Trade-Offs, and IPv6 Transition Challenges

Yousra Almukhtar*[1], Nuredin Ahmed[2]

[1]*Department of Information Technology, Libya Academy for Graduate Studies, Tripoli, Libya*
[2]*Department of Computer Engineering, University of Tripoli, Tripoli, Libya*
Corresponding email. *Yosraalmokhtar4@gmail.com*

Abstract
For over a decade, the global internet has navigated a depleted IPv4 address pool. The telecommunications industry has maintained legacy operations through complex mitigation frameworks, primarily Dual-Stack, Tunneling, and stateful Translation mechanisms (e.g., CGNAT and PAT). To evaluate the long-term architectural and economic implications of these strategies, this review employs a thematic synthesis of 24 peer-reviewed studies and recent empirical simulations published over the last ten years, sourced from major academic databases using specific inclusion criteria. The thematic synthesis applies a structured coding framework to extract performance (latency, packet loss), security, and economic indicators across the selected studies, enabling comparative evaluation. The analysis reveals a structural paradox: while these mechanisms successfully bypass immediate address exhaustion and scale existing networks, they often compromise the Internet's end-to-end design and introduce complex security vulnerabilities. The findings indicate that stateful translation mechanisms introduce measurable latency increases (up to reported 15–30% in high-load scenarios) and scalability bottlenecks, while Dual-Stack architectures significantly increase resource utilization. Furthermore, the persistence of mitigation frameworks contributes to the expansion of the IPv4 leasing market, creating economic disincentives for IPv6 adoption. This review is limited by reliance on publicly available datasets and simulation-based studies, with restricted access to proprietary ISP operational metrics. Ultimately, current transition mechanisms serve as necessary pragmatic interim phases rather than sustainable permanent solutions, highlighting the critical need for a definitive shift toward native IPv6.
Keywords: IPv4 Exhaustion; IPv6 Transition; Dual-Stack; CGNAT; IPv4aaS; Network Scalability; Architectural Trade-Offs.

Introduction

The Internet's foundational architecture relies on IPv4, a protocol that successfully supported early network demands [1], [2]. However, the rapid expansion of mobile and cloud computing has exhausted this 32-bit addressing space [3]. Although global IPv6 adoption reached an average of 46% by 2025—exceeding 74% in leading regions [4], the complete deprecation of IPv4 remains delayed. Consequently, the depletion of unallocated pools across major registries has transformed IPv4 addresses into scarce economic assets, with reported secondary-market prices averaging approximately $20–25 per IPv4 address in recent studies [5], [6], subject to regional variation, which continues to impose heavy operational expenditures (OPEX).

To maintain connectivity, the industry relies on interim mitigation strategies, including Dual-Stack environments [7], [8] and stateful translation mechanisms like Carrier-Grade NAT (CGNAT) and PAT [9], [10]. While empirical models show these sustain legacy operations, a critical research gap remains: current literature lacks a formal synthesis evaluating their combined quantitative and architectural toll on core infrastructure. This study addresses the following research questions:

- **(RQ1)** What are the performance, security, and economic trade-offs of current IPv4 mitigation techniques?
- **(RQ2)** To what extent do these techniques delay or facilitate IPv6 adoption?
- **(RQ3)** What architectural limitations constrain their long-term sustainability?

Addressing this gap, this paper systematically reviews these strategies to analyze their structural trade-offs, including increased routing overhead [11] and expanded security vulnerabilities [12]. Ultimately, this study argues that the architectural and economic debt of prolonged IPv4 mitigation is unsustainable, emphasizing the long-term operational constraints that necessitate an accelerated shift to native IPv6.

Methodology

Guided by the PRISMA framework to ensure reproducibility, this review employs a systematic thematic synthesis.

Search Strategy & Selection Process

A targeted search across IEEE Xplore, ACM Digital Library, and Google Scholar covering the last ten years utilized Boolean strings: ("IPv4 exhaustion" OR "IPv4 depletion") AND ("CGNAT" OR "Dual-Stack" OR "IPv6 transition"). The literature search was conducted between January 2025 and March 2026. Initial query results yielded 35 studies (IEEE Xplore: 14, ACM: 9, Google Scholar: 12). Literature was rigorously filtered:

Received*: 05-04-2026 -* ***Accepted****: 02-06-2026 -* ***Published****: 09-06-2026*

- **Inclusion:** Peer-reviewed studies providing empirical metrics, economic scalability data, or formal security analysis (e.g., STRIDE) [12].
- **Exclusion:** Descriptive tutorials, non-peer-reviewed whitepapers, and deprecated mechanisms (e.g., 6to4).

Of 35 initially identified papers, 11 were excluded for lacking quantitative rigor or direct architectural relevance, finalizing a corpus of 24 high-impact studies. The selection process follows PRISMA guidelines, with screening stages including identification, eligibility assessment, and final inclusion.

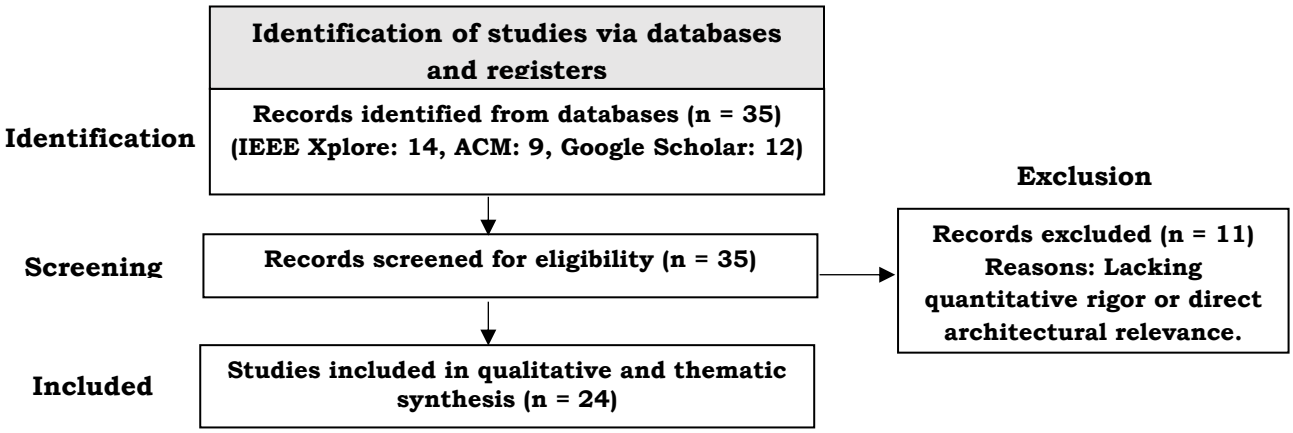

Fig. 1: PRISMA flow diagram detailing the study selection process

Study Quality Assessment

Each study was evaluated using three criteria: methodological rigor, relevance to architectural evaluation, and presence of quantitative metrics (e.g., latency, throughput, or economic cost). Studies lacking empirical validation were excluded.

Study Classification & Thematic Synthesis The 24 selected studies were systematically analyzed. Themes were identified through iterative coding of extracted data and grouped into four analytical domains: performance, security, economic scalability, and long-term viability. This analytical matrix classified the transition mechanisms into three primary frameworks:

Co-existence (Dual-Stack Architecture)

This strategy requires network nodes to run both IPv4 and IPv6 protocols simultaneously [7], [8]. Empirical simulations demonstrate that while Dual-Stack environments ensure seamless compatibility and maintain native routing efficiency without encapsulation overhead [8], they introduce distinct operational trade-offs. The architecture increases hardware resource utilization and administrative complexity without mitigating the underlying IPv4 scarcity [13]. Reported studies indicate CPU utilization increases ranging between 10% and 25% when maintaining dual routing tables in high-throughput environments [8], [13]. Implementing this framework across a hierarchical topology requires every layer—from the core routing matrix to edge access devices—to actively maintain dual routing tables and dynamic configurations like OSPFv3 [8]. Consequently, this structural duplication remains a primary driver of the elevated operational and configuration expenditures associated with this transitional phase. As illustrated in Fig.2, this dual protocol processing leads to duplication of routing and forwarding logic across network layers, directly contributing to increased hardware utilization.

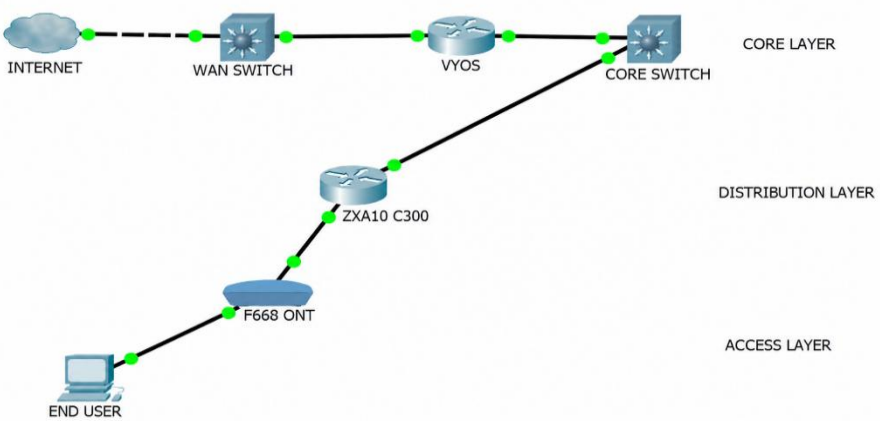

Fig. 2. Architectural Layout of an ISP Dual-Stack Laboratory Network Environment. (Adapted from [16])

Received*: 05-04-2026 -* ***Accepted****: 02-06-2026 -* ***Published****: 09-06-2026*

Tunneling Mechanisms

Unlike the native routing efficiency of Dual-Stack architectures [8], this framework encapsulates IPv6 packets within IPv4 headers to traverse legacy networks [14]. While it facilitates necessary interoperability across disjointed infrastructures, the encapsulation process introduces substantial protocol overhead. Consequently, rather than violating theoretical ideals, this added complexity tangibly increases latency, causes Maximum Transmission Unit (MTU) mismatches, and creates potential points of failure, presenting distinct performance limitations compared to direct routing models [11].

Translation and Address Sharing Mechanisms

Technologies such as CGNAT and 464XLAT map multiple private hosts to a single public IPv4 address using port-based translation (PAT) [10], [15], [16]. Recent empirical simulations confirm that while these techniques efficiently scale networks and provide immediate relief to address depletion [10], they introduce significant structural trade-offs. Primarily, they disrupt the Internet's native end-to-end communication model by forcing resource-heavy, stateful processing onto network routers. As demonstrated in Fig. 3, this architecture typically manifests through two successive NAPT44 translation layers executed sequentially at the Customer Premises Equipment (CPE) and the ISP gateway. This architecture highlights the compounding effect of sequential NAT layers, which increases latency and complicates traceability. This dual-stage framework requires network nodes to continuously track and maintain exhaustive stateful translation tables for every active session. Consequently, while successfully extending IPv4 utility, this continuous state-tracking introduces latency increases commonly reported between 5–30 ms under high concurrent session loads, alongside measurable packet loss in state-table saturation scenarios [10], [16], [18], and leaves core gateways highly susceptible to state-exhaustion vulnerabilities. However, some studies argue that optimized CGNAT deployments with hardware acceleration can mitigate performance penalties, suggesting that implementation quality significantly affects observed outcomes.

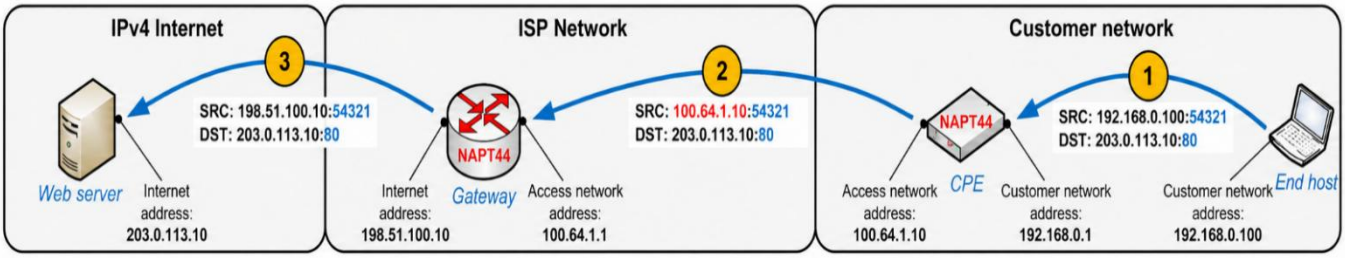

Fig.3. Architecture of Translation and Address Sharing Mechanisms (Dual-NAPT44/CGNAT). (Adapted from [15])

Results and Discussion

Architectural Critique and Research Gaps:

This section critically synthesizes empirical findings to expose the architectural trade-offs of current IPv4 mitigation strategies. By confronting the literature's contradictions, it identifies a core research gap: the growing disconnect between short-term engineering fixes and long-term structural realities.

The Co-existence Model: Strengths vs. Architectural Weaknesses

While the Dual-Stack strategy successfully preserves native protocol integrity and backward compatibility [7], [8], recent empirical simulations highlight its substantial infrastructure overhead. Operating parallel protocols, such as concurrent OSPFv3 and IPv4 routing, significantly increases CPU utilization and memory consumption within core routers [8], [13], [20]. Reported studies indicate CPU utilization increases ranging between 10% and 25% when maintaining dual routing tables in high-throughput environments [8], [13]. Furthermore, this architectural duplication inherently expands the network's attack surface. Managing dual firewalls significantly increases administrative complexity and configuration overhead—a practical vulnerability recently acknowledged in modern transition models that frequently defer advanced IPv6 security configurations due to their operational difficulty [8], [11]. Consequently, while Dual-Stack ensures transitional stability, its hardware and security demands present challenging long-term trade-offs.

Translation Complexity and the Core Research Gap

Translation mechanisms, including CGNAT and PAT, have successfully extended IPv4's operational lifespan, providing crucial immediate relief and scalability for ISPs [10], [17],[21]. However, recent simulations and empirical benchmarks demonstrate that this extension incurs substantial performance costs. Under high concurrent loads, the continuous state-tracking required by these mechanisms introduces latency increases commonly reported between 5–30 ms under high concurrent session loads, alongside

Received*: 05-04-2026 -* ***Accepted****: 02-06-2026 -* ***Published****: 09-06-2026*

measurable packet loss in state-table saturation scenarios [10], [16], [18]. However, some studies argue that optimized CGNAT deployments with hardware acceleration can mitigate performance penalties, suggesting that implementation quality significantly affects observed outcomes.

This technical trade-off highlights a broader empirical disconnect. Rather than accelerating native IPv6 migration, the industry's reliance on these interim solutions has inadvertently fueled a lucrative, speculative IPv4 leasing market [5], [6],[22]. Consequently, networks have adopted complex, multi-layered hybrid frameworks (e.g., Class 2 architectures). As illustrated in Fig. 4, mitigating address scarcity through these models simply transfers the architectural burden to the core network. Sequential tunneling and NAPT44 translation demand excessive state tracking at the centralized ISP gateway, significantly increasing protocol overhead and establishing an operationally challenging compromise rather than a sustainable resolution.

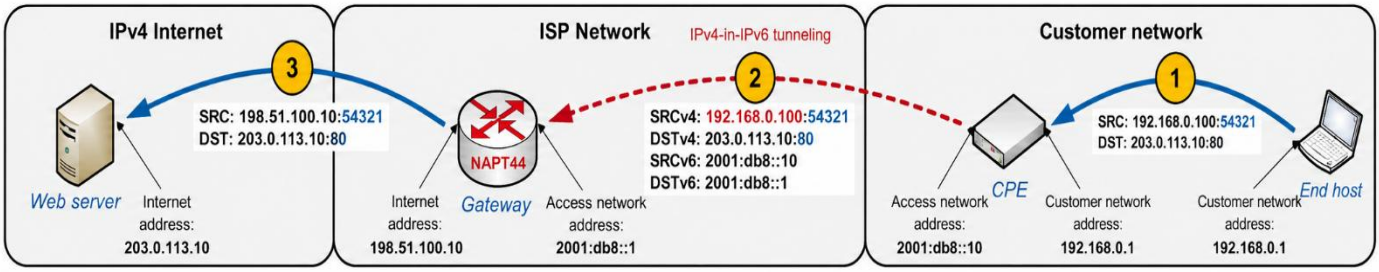

Fig. 4. Escalated architectural complexity in a Class 2 hybrid sharing mechanism. (Adapted from [15])

Comprehensive Evaluation Matrix

Table 1 systematically evaluates IPv4 mitigation strategies using an objective rating scale derived from empirical benchmarks. 'Performance' is assessed via latency, packet loss, and encapsulation overhead metrics [8], [10], [16], [18]; 'Security' by attack surface expansion and traceability [11], [12]; 'Economic Scalability' by CAPEX/OPEX and IPv4 leasing dependencies [5], [6],[23]; and 'Long-Term Viability' by the capacity to facilitate native IPv6 migration [9], [17], [24].

Table 1: Critical Evaluation Matrix of IPv4 Mitigation Strategies

Mitigation Strategy	Performance & Overhead	Security Implications	Economic Scalability	Long-Term Viability
Dual-Stack [7], [8], [13]	Native routing; zero encapsulation overhead [8].	Doubled attack surface; complex synchronized firewalls [8], [11].	High memory (TCAM) costs for dual BGP tables.	Vital interim step; delays migration; does not solve exhaustion.
Tunneling [14], [19]	Fragmentation; MTU mismatches; processing delays.	Vulnerable to spoofing and injection; bypasses perimeter defenses.	Low CAPEX; high unpredictable OPEX; complex troubleshooting.	Unstable for modern topologies; largely deprecated.
Stateful Translation (e.g., CGNAT, PAT) [10], [15], [17]	High latency; state-table exhaustion; packet loss under load [10], [16].	Loss of end-to-end IP attribution; crippled abuse tracking.	Expensive carrier-grade hardware; costly IPv4 leasing [5].	Architectural bottleneck; disincentivizes IPv6 adoption [9].
Stateless / Hybrid (e.g., 464XLAT) [16]	Eliminates state bottlenecks; minor encapsulation overhead.	Improved traceability; requires intermediate policy sync [12].	Highly scalable; offloads state processing to CPE.	Viable interim IPv4aaS model during IPv6-only core migration.

Ratings are derived from synthesized findings across multiple empirical and simulation-based studies.

Future Directions

To address the systemic gaps identified in this review, future research must abandon the optimization of legacy workarounds and pivot toward three critical domains:

- **Bridging the Architectural Gap:** Shift focus from complex hybrid transition mechanisms to the operationalization of native IPv6-only architectures, developing frameworks to gracefully phase out resource-heavy Dual-Stack environments.
- **Bridging the Economic Gap:** Propose interdisciplinary regulatory policies and economic models designed to systematically depreciate the speculative value of legacy IPv4 blocks, thereby financially compelling market migration.

Received*: 05-04-2026 -* ***Accepted****: 02-06-2026 -* ***Published****: 09-06-2026*

- **Bridging the Security Gap:** Design robust, native IPv6 defense frameworks that restore end-to-end IP traceability and accountability, eliminating the reliance on stateful NAT layers for perimeter security.
-

Conclusion

In conclusion, this critical review demonstrates that while IPv4 mitigation strategies—such as Dual-Stack and stateful translation—prevent immediate address exhaustion, they introduce significant architectural overhead and inadvertently fuel an IPv4 leasing market that disincentivizes native IPv6 adoption.
However, an immediate shift to pure IPv6 is hindered by profound real-world constraints. High capital expenditures (CAPEX) and complex legacy integrations make sudden migration unfeasible, rendering current transition mechanisms a pragmatic and necessary interim phase rather than an architectural failure. From a practical perspective, network operators should prioritize phased IPv6-only core deployment while minimizing reliance on stateful translation. Policymakers may also consider regulatory mechanisms to reduce speculative IPv4 market dynamics. Finally, these findings should be interpreted within the study's limitations. Bounded by a synthesis of 24 peer-reviewed studies from the past decade, this review may not capture undisclosed proprietary data from global ISPs. Ultimately, while mitigation strategies are essential for current stability, long-term operational sustainability dictates a strategically phased but inevitable transition toward IPv6-only network architectures.

References

1. Muni Y. Understanding the evolution of internet protocols: an in-depth review of IPv4 and IPv6: a comparative review of transition challenges and solutions. Int J Adv Res Comput Sci. 2025;16(4):109-17.
2. Bakni M, Hanbo S. A survey on Internet Protocol version 4 (IPv4). WikiJournal Sci. 2022;5(1):2.
3. Hossain MM, Binti J, Uddin MK. A review paper on IPv4 and IPv6: a comprehensive survey. Am J Comput Sci Technol. 2024;7(4):170-5.
4. Ugwumba NK. Empirical analysis of IPv6 protocol implementation: a comprehensive framework for address validation, network discovery, and security assessment. Res Square [Preprint]. 2025.
5. Degen B, et al. From scarcity to opportunity: examining abuse of the IPv4 leasing market. Proc 9th Netw Traffic Meas Anal Conf (TMA). 2025:1-11.
6. Prehn L, Lichtblau F, Feldmann A. When wells run dry: the 2020 IPv4 address market. Proc 16th Int Conf Emerging Netw Experiments Technol. 2020:46-54.
7. Amarnath VK. IPv6 adoption alongside IPv4: a strategic dual-stack approach for network sustainability. J Inf Syst Eng Manag. 2025;10(61s):977-87.
8. ALgzite H, Ahmed N. IPv6 addressing and configuration: building a dual-stack network with OSPFv3. Alqalam J Med Appl Sci. 2025;8(4):2974-80.
9. Salamatian L, et al. Who squats IPv4 addresses? ACM SIGCOMM Comput Commun Rev. 2023;53(1):48-72.
10. Abudaber T, Ahmed N. Implementation and analysis of NAT and PAT techniques in Cisco-based networks for efficient IPv4. Alqalam J Med Appl Sci. 2026;9(1):58-63.
11. Ashraf Z, Sohail A, Latif S, Hameed A, Yousaf M. Challenges and mitigation strategies for transition from IPv4 network to virtualized next-generation IPv6 network. Int Arab J Inf Technol. 2023;20(1).
12. Al-Azzawi A, Lencse G. Methodology for the security analysis of IPv4-as-a-Service IPv6 transition technologies. Comput J. 2025;68(10):1450-62.
13. Paul HC, Bakon KA. A study on IPv4 and IPv6: the importance of their co-existence. Int J Inf Syst Eng. 2016;4(2):97-106.
14. Ashraf S, Muhammad D, Aslam Z. Analyzing challenging aspects of IPv6 over IPv4. J Ilm Tek Elektro Komput Dan Inform. 2020;6(1):54-67.
15. Škoberne N, Maennel O, Phillips I, Bush R, Žorž J, Ciglarič M. IPv4 address sharing mechanism classification and trade-off analysis. IEEE/ACM Trans Netw. 2014;22(1):182-95.
16. Lencse G, Nagy N. Towards the scalability comparison of the Jool implementation of the 464XLAT and of the MAP-T IPv4aaS technologies. Int J Commun Syst. 2022;35(17):e5354.
17. Hamdou M, El Ksimi A, Ettaki B. Classification of IPv6 transition mechanisms using multiple-criteria decision-making. Eng Technol Appl Sci Res. 2025;15(3):22960-8.
18. Hamarsheh A, et al. Comparative evaluation of host-based translator mechanisms for IPv4-IPv6 communication performance analysis with different routing protocols. Int J Cloud Appl Comput. 2023;13(1):1-26.
19. Cordeiro E, Carnier R, Zucchi WL. Comparison between IPv4 to IPv6 transition techniques. arXiv preprint arXiv:1612.01948. 2016.
20. Cañas R, Henríquez-Miranda C, Silva J. Comparative analysis of IPv4 and IPv6 to improve quality of service on a university wireless network. CESTA. 2025;13(14):1-6.
21. Enache D, Alexandru M. A study of the technology transition from IPv4 to IPv6 for an ISP. Rev Air Force Acad. 2016;14(1):117-22.
22. Hamarsheh A, Abdalaziz YP, Nashwan S. Recent impediments in deploying IPv6. Adv Sci Technol Eng Syst J. 2021;6(1):336-41.
23. Kodakandla NN. IPv4 vs. IPv6 in cloud engineering: performance, security and cost analysis. Int J Sci Res Arch. 2023;8(2):774-84.
24. Sun BS. IPswEN: a long term evolution approach for the IPv4 Internet architecture. Proc IFIP Netw Conf. 2022:1-9.

Received*: 05-04-2026 -* ***Accepted****: 02-06-2026 -* ***Published****: 09-06-2026*