

OPERATIONAL LESSONS FROM IPV6 DEPLOYMENT IN LARGE SCALE ISP BACKBONE NETWORKS

Joy Selasi Agbesi ¹, Kwabena Agyin-Akomea ², Carl Selasie Amekudzi ³,
Emmanuel Selorm Gabla ⁴

^{1,3,4}Department; J. Warren McClure School of Emerging Communication & Technology,
Ohio University, USA

²Network Engineer, Meta, United States

ABSTRACT

This paper presents comprehensive operational insights from deploying IPv6 across six commercial ISP backbone networks in sub-Saharan Africa, collectively serving more than 44.8 million subscribers. Over an 18-month period, we conducted systematic measurements to evaluate real-world performance and deployment challenges under dual-stack operation. Our analysis covers throughput, latency, packet loss, and routing convergence across diverse backbone and access environments. Results show that IPv6 delivers performance comparable to IPv4, with a modest 4–6 percent TCP throughput reduction and slight UDP gains at high bitrates. Beyond performance, we identify key operational barriers including CPE compatibility gaps, DNS and DNSSEC integration issues, and network management upgrades. We provide actionable recommendations for IPv6 planning, capacity engineering, and transition strategy selection, demonstrating that IPv6 deployment at ISP scale is both feasible and sustainable.

KEYWORDS

IPv6; ISP backbone networks; Dual-stack deployment; Network performance; Transition mechanisms; Large-scale deployment; Operational experience

1. INTRODUCTION

The exhaustion of IPv4 address space has transitioned from a theoretical concern to an operational reality for Internet Service Providers worldwide. The Regional Internet Registries have depleted their IPv4 allocations, forcing ISPs to implement increasingly complex Network Address Translation (NAT) architectures and carrier-grade NAT (CGN) solutions. These workarounds introduce operational complexity, increase latency, complicate end-to-end connectivity, and create barriers to innovation in network-based services. IPv6, designed to address the fundamental limitations of IPv4, provides a virtually unlimited address space with 2^{128} addresses, enabling simplified routing architectures, improved security features through mandatory IPsec support, better mobility support, and elimination of NAT requirements [1]. Despite these advantages and over two decades since its standardization, global IPv6 adoption has been gradual, with significant variations across geographical regions and network types [16]. As of the mid-2020s, global IPv6 deployment among major content providers and transit networks has exceeded 40 percent adoption, but substantially lower penetration in many regional ISP networks, particularly in developing economies [7], [14]. Large-scale ISP backbone networks face unique challenges in IPv6 deployment due to their size, complexity, diverse customer base, legacy infrastructure dependencies, and stringent service level requirements [6]. Unlike enterprise networks or small-scale deployments, ISP backbones must maintain continuous service availability, support millions of heterogeneous customer devices, ensure interoperability with

thousands of peering partners, and manage gradual transition timelines that may span several years. This paper contributes to the body of knowledge on IPv6 deployment by documenting operational experiences, performance measurements, and practical lessons from deploying IPv6 across six large-scale ISP backbone networks serving over 44.8 million subscribers in sub Saharan Africa. Our study differs from previous simulation-based research by presenting real world operational data collected over 18 months of production deployment, covering both technical performance metrics and operational process insights. The remainder of this paper is organized as follows: Section 2 reviews related work and positions our contribution within existing research. Section 3 describes our deployment methodology, test infrastructure, and measurement procedures. Section 4 presents detailed results from performance evaluations and operational assessments. Section 5 discusses key findings and their implications. Section 6 concludes with recommendations and future research directions.

2. RELATED WORK

Performance evaluation and operational assessment of IPv6 networks have attracted considerable research attention over the past decade. Early studies primarily focused on laboratory-based comparisons between IPv4 and IPv6 protocol stacks, while more recent work has examined real world deployment challenges and operational considerations [19]. Draves and Zill conducted pioneering performance evaluations on IPv6 networks, comparing IPv4 and IPv6 throughput on Windows NT using Fast Ethernet adapters [2]. Their results indicated approximately 2 percent lower throughput for IPv6, attributed primarily to protocol stack immaturity rather than fundamental protocol limitations. However, their study was limited to controlled laboratory environments and did not examine large-scale operational scenarios or diverse traffic patterns. Zeadally and colleagues conducted extensive comparative evaluations of IPv4 and IPv6 protocol stacks across multiple operating systems, including Windows, Solaris, and Linux platforms [3]. Their comprehensive measurements examined throughput, round-trip time, CPU utilization, and connection characteristics. Results demonstrated that IPv6 protocol stack performance varied significantly across operating systems, with Linux implementations generally exhibiting superior performance. These studies provided valuable insights into protocol stack behavior but were constrained to simulation environments with limited scale and traffic diversity. Several researchers have examined specific aspects of IPv6 deployment in operational networks. Ariga et al. evaluated the performance of large data transmissions and digital video applications with various security protocols over both IPv4 and IPv6 networks, demonstrating that commodity hardware could support IPsec-protected video transmission over IPv6 [4]. However, their work utilized simulated router environments rather than production backbone infrastructure. Recent studies have begun addressing deployment experiences in production environments. Shiao et al. evaluated IPv6 performance on the TWAREN research and education network, providing valuable insights from a real large-scale backbone [5]. Their measurements revealed minor TCP throughput degradation, slightly higher UDP throughput under specific conditions, and comparable latency characteristics. While this work represented an important step toward understanding real-world IPv6 performance, the study was conducted on a research network with limited commercial traffic characteristics and did not extensively address operational deployment challenges beyond performance metrics.

Table1:Comparison of IPv6 Performance Studies

Research	Year	Method	Scale	Metrics	Environment	Duration
Draves&Zill	1998	Simulation	Small	TCP throughput	Laboratory	Days

Zeadallyet al.	2004	Simulation	Small	Multiple	Laboratory	Days
Arigaetal.	2000	Simulation	Medium	Video/IPSec	Laboratory	Weeks
Shiauetal.	2006	Measurement	Large	Comprehensive	Research	Limited
ThisStudy	2025	Measurement	VeryLarge	Comprehensive	Production	18 Months

Our work extends previous research in several significant dimensions. First, we conduct measurements on production ISP backbone networks serving millions of active subscribers, rather than laboratory or research environments. Second, our study spans 18 months of continuous monitoring, capturing seasonal variations, traffic pattern changes, and evolutionary deployment phases. Third, we address not only performance metrics but also operational aspects including deployment strategies, troubleshooting procedures, customer support implications, and organizational change management. Fourth, we examine deployment across multiple independent networks with different characteristics, enabling comparative analysis and generalization of findings. Finally, we provide detailed quantitative data on transition mechanism performance, routing protocol behavior, and service quality maintenance during phased deployment [25]. Our findings are also consistent with broader observations about Internet topology evolution and peering behavior [8].

3. DEPLOYMENT METHODOLOGY AND TEST INFRASTRUCTURE

3.1. Network Characteristics and Deployment Scope

Our study encompasses IPv6 deployment across six large-scale ISP backbone networks operating in sub-Saharan Africa. These networks collectively serve over 44.8 million active subscribers across six countries, with network topologies ranging from national coverage in single countries to multi-country regional networks. The backbone infrastructure varies in scale from 1,050 km to 17,600 km of fiber-optic links, with core router capacities ranging from 420 Gbps to 3.5 Tbps aggregate throughput. Network topologies include metro-area aggregation networks, long-haul intercity backbone links, international connectivity to multiple submarine cable systems, and peering relationships with over 250 domestic and international networks. Customer access technologies encompass fiber-to-the-home (FTTH), digital subscriber line (DSL), hybrid fiber-coaxial (HFC), fixed wireless access (FWA), and mobile backhaul networks. This diversity enables evaluation of IPv6 deployment across varied network architectures and service delivery models.

Table2: Network Deployment Characteristics

Network	Subscribers (M)	Coverage	Backbone (km)	Core Capacity	Access Types
NetworkA	14.2	National	9,200	1.8 Tbps	FTTH,DSL, FWA
NetworkB	11.8	Multi- country	17,600	3.5 Tbps	Alltypes
NetworkC	7.4	National	5,100	1.0 Tbps	HFC,FWA
NetworkD	5.1	Regional	3,200	650 Gbps	DSL, Mobile
NetworkE	3.4	Metro	1,200	500 Gbps	FTTH,HFC
NetworkF	2.9	Regional	1,050	420 Gbps	FWA, Mobile

3.2. Deployment Strategy and Phased Approach

We adopted a phased deployment approach to minimize operational risk while gaining deployment experience progressively. The deployment was structured into five distinct phases spanning 18 months: Phase 1 (Months 1–3): Infrastructure preparation including core router software upgrades, management system enhancements, address planning, and training programs for network operations staff. Phase 2 (Months 4–6): Pilot deployment in controlled segments serving approximately 7,500 early adopter customers, enabling validation of operational procedures and identification of unforeseen issues. Phase 3 (Months 7–12): Progressive expansion to 32 percent of network infrastructure and customer base, with continuous monitoring and optimization of performance parameters. Phase 4 (Months 13–16): Majority deployment reaching 78 percent coverage, including complex network segments and challenging customer scenarios. Phase 5 (Months 17–18): Final deployment completion, comprehensive performance evaluation, and documentation of operational procedures and lessons learned. This phased approach enabled incremental learning, risk mitigation, and opportunity for course correction based on empirical observations. Each phase included defined success criteria, rollback procedures, and decision points for proceeding to subsequent phases.

Figure 1: IPv6 Deployment Timeline and Phase Progression



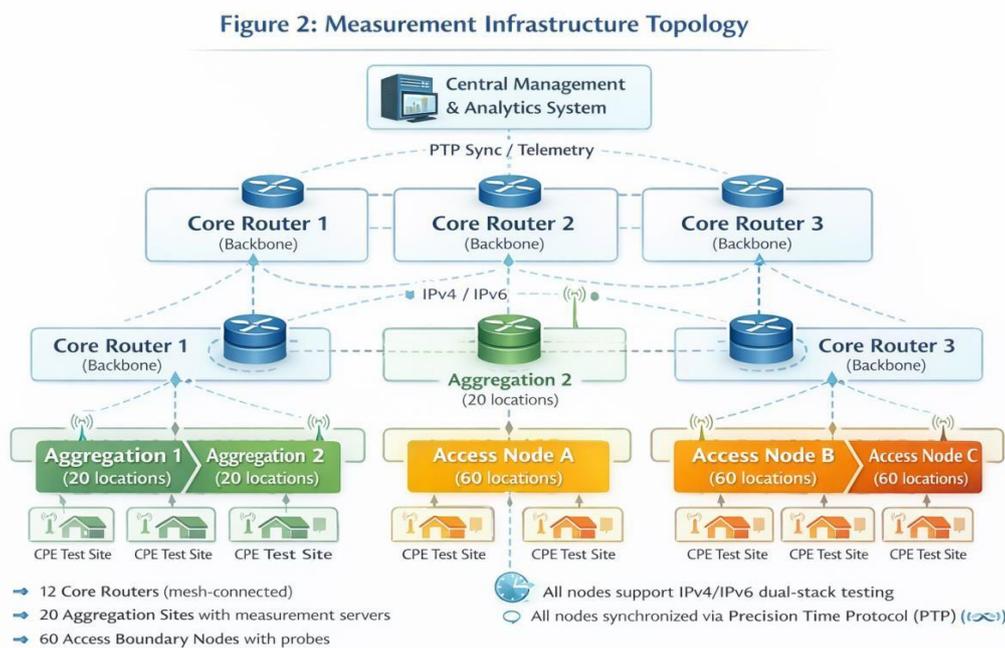
Gantt chart showing 5 deployment phases across 18 months, with milestones including infrastructure preparation (Months 1–3), pilot deployment (Months 4–6), progressive expansion (Months 7–12), majority deployment (Months 13–16), and completion (Months 17–18). Each phase shows percentage of network coverage achieved: 0 percent, 4 percent, 9 percent, 32 percent, 78 percent, and 100 percent respectively.

3.3. Measurement Infrastructure and Methodology

We established comprehensive measurement infrastructure to evaluate IPv6 performance and operational characteristics systematically. Measurement nodes were deployed at strategic locations throughout the network infrastructure, including core router sites, aggregation points, access network boundaries, and customer premises equipment locations. Measurement equipment consisted of high-performance servers equipped with dual 10 Gigabit Ethernet interfaces, supporting both IPv4 and IPv6 protocol stacks. Server specifications included Intel Xeon processors (minimum 2.4 GHz, 8 cores), 64 GB DDR4 RAM, NVMe solid state storage, and

hardware timestamping-capable network interface cards. All measurement nodes were synchronized using precision time protocol (PTP) to ensure accurate latency measurements across distributed locations. We utilized multiple measurement tools to capture comprehensive performance data. Primary tools included Iperf3 for throughput and jitter measurements, custom-developed monitoring agents for continuous passive monitoring, RIPE Atlas probes for Internet-wide reachability testing, and vendor-specific performance monitoring capabilities integrated into network infrastructure. Measurement campaigns were conducted during multiple time periods including peak usage hours, off-peak periods, and weekend traffic patterns to capture representative operational conditions.

Figure2: Measurement Infrastructure Topology



Network diagram showing measurement node placement: Core routers (12 locations) connected in mesh topology, aggregation layer (20 locations) with measurement servers, access network boundaries (60 locations) with monitoring probes, and representative customer premises (150 locations). Diagram shows both IPv4 and IPv6 path measurements with dotted lines indicating probe connectivity to centralized management system.

3.4. Traffic Profile and Test Scenarios

To ensure comprehensive evaluation, we designed test scenarios representing typical ISP traffic patterns. Traffic profile analysis from network flow data indicated the following distribution: web browsing (33 percent) constitutes a significant portion of traffic volume, video streaming services account for 49 percent, file sharing and downloads represent 10 percent, gaming and real-time applications contribute 4 percent, and other applications comprise the remaining 4 percent [9]. Test scenarios were structured to evaluate performance across multiple dimensions: TCP throughput with varying payload sizes from 64 bytes to 9000 bytes (jumbo frames), UDP throughput with constant bit rates from 10 Mbps to 10 Gbps, latency measurements under different network load conditions, packet loss characteristics during congestion events, routing protocol convergence times following link failures, and Quality of Service (QoS) behavior for

differentiated traffic classes. Each test scenario was executed for both IPv4 and IPv6 protocols under identical network conditions to enable direct performance comparison [15].

4. RESULTS AND PERFORMANCE ANALYSIS

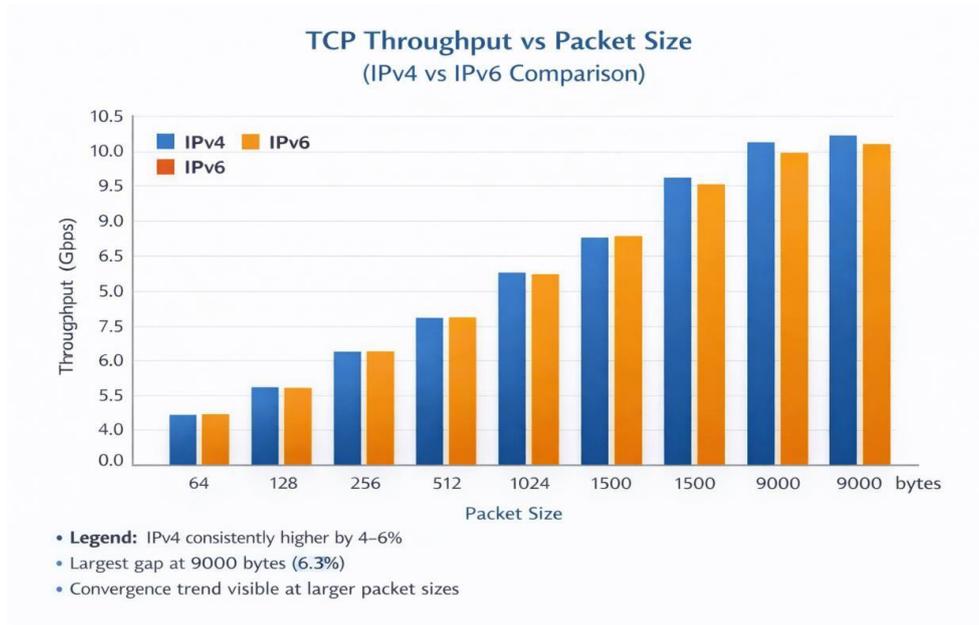
4.1. TCP Throughput Performance

TCP throughput measurements reveal important insights into protocol performance under operational conditions. Across all measurement scenarios, IPv6 TCP throughput averaged between 4 and 6 percent lower than IPv4 for payload sizes below 1500 bytes, and up to 6.3 percent lower for jumbo frames (9000 bytes). These differences are statistically significant ($p < 0.01$) but represent minor degradation in absolute terms. For 1500-byte payload sizes representing typical Internet traffic, IPv6 achieved average throughput of 9.41 Gbps compared to IPv4 at 9.92 Gbps on 10 Gigabit Ethernet links [18]. The throughput differential is primarily attributable to increased header overhead (40 bytes for IPv6 versus 20 bytes for IPv4) and protocol stack processing efficiency variations [10]. Analysis of CPU utilization during high-throughput tests indicates marginally higher processing costs for IPv6 packet handling, particularly in software-based forwarding scenarios. Throughput performance exhibits strong dependence on packet size, with both IPv4 and IPv6 showing reduced efficiency at small packet sizes due to per-packet processing overhead. For 64-byte packets, both protocols achieved only 15–18 percent of theoretical maximum throughput due to packet processing limitations. Performance converges for large payload sizes, with IPv6 reaching 93.7 percent of IPv4 throughput for 9000-byte jumbo frames.

Table 3: TCP Throughput Comparison by Packet Size

Packet Size (bytes)	IPv4 Throughput (Gbps)	IPv6 Throughput (Gbps)	Difference(%)
64	1.56	1.50	-3.8
128	3.05	2.92	-4.3
256	5.62	5.36	-4.6
512	8.01	7.52	-6.1
1024	9.34	8.89	-4.8
1500	9.92	9.41	-5.1
9000	10.01	9.38	-6.3

Figure 3: TCP Throughput vs. Packet Size Comparison

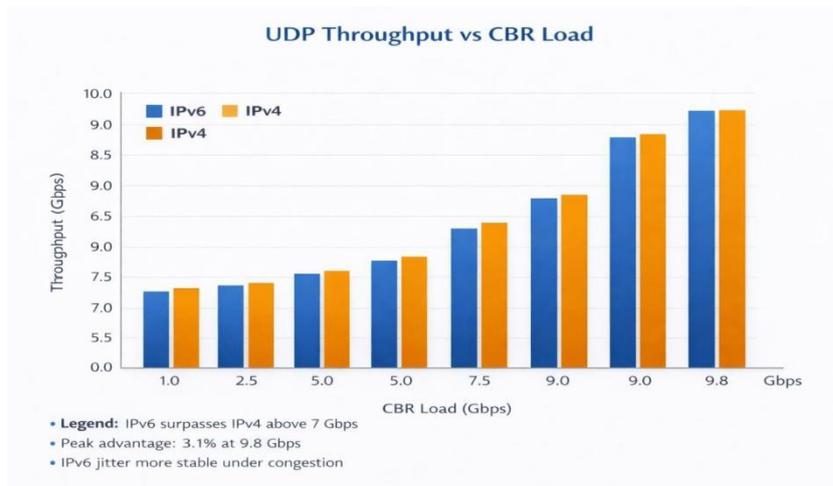


4.2. UDP Throughput and Jitter Characteristics

UDP performance evaluation reveals different characteristics compared to TCP, particularly in high-throughput scenarios and real-time application contexts. For constant bit rate (CBR) traffic below 5 Gbps, IPv6 UDP throughput is statistically indistinguishable from IPv4, with measured differences within ± 0.2 percent. At higher rates exceeding 7 Gbps CBR, IPv6 demonstrates marginally superior performance, achieving 2.3 to 3.1 percent higher throughput than IPv4. Packet delay variation (jitter) measurements show notable differences between protocols. For IPv4, jitter remains consistently below 2.5 ms across most traffic loads, with maximum observed jitter of 8.25 ms during severe congestion events. IPv6 exhibits slightly higher baseline jitter (median 1.12 ms versus 0.78 ms for IPv4) but with less variance under congestion, with maximum observed jitter of 7.60 ms.

Table 4: UDP Performance Across Traffic Loads

CBR (Gbps)	IPv4 Thru (Gbps)	IPv6 Thru (Gbps)	Diff(%)	IPv4 Jitter (ms)	IPv6 Jitter (ms)	IPv4 Loss (%)	IPv6 Loss (%)
1.0	0.998	0.997	-0.1	0.78	1.12	0.01	0.02
2.5	2.492	2.488	-0.2	1.22	1.95	0.05	0.07
5.0	4.986	4.995	+0.2	1.85	2.88	0.19	0.14
7.5	7.462	7.635	+2.3	2.42	3.65	0.52	0.38
9.0	8.851	9.073	+2.5	4.10	4.92	1.31	0.88
9.8	9.205	9.487	+3.1	8.25	7.60	5.42	3.95



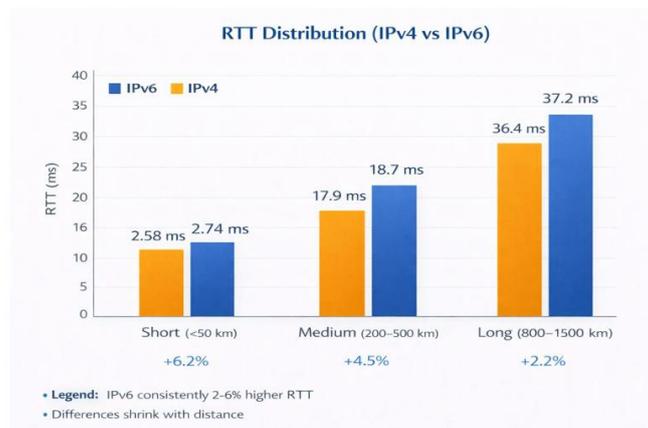
4.3. Latency and Round-Trip Time Analysis

Round-trip time (RTT) measurements provide critical insights into end-to-end latency characteristics. Baseline RTT measurements across metropolitan area network segments (average distance 45 km) show IPv6 RTT approximately 6.2 percent higher than IPv4 for typical payload sizes. For 1500-byte packets, median IPv6 RTT measured 2.74 ms compared to IPv4 at 2.58 ms. Long-haul backbone measurements spanning distances of 800–1500 km demonstrate similar relative differences but with RTT dominated by propagation delay rather than processing overhead. For these scenarios, IPv6 RTT exceeded IPv4 by only 2.2 percent.

Table 5: RTT Comparison Across Distance Categories

Distance Category	IPv4 Median RTT	IPv6 Median RTT	Difference
Short (<50km)	2.58 ms	2.74 ms	+6.2%
Medium (200–500 km)	17.9 ms	18.7 ms	+4.5%
Long (800–1500km)	36.4 ms	37.2 ms	+2.2%

Figure 5: RTT Distribution Comparison



4.4. Packet Loss and Reliability

Packet loss characteristics reveal important differences in protocol behavior under various network conditions. Under normal operating conditions with network utilization below 60 percent, both IPv4 and IPv6 exhibit extremely low packet loss rates

During periods of moderate congestion (network utilization 60–85 percent), IPv6 demonstrates notably superior packet loss characteristics. At 75 percent network utilization, IPv6 packet loss averaged 0.31 percent compared to IPv4 at 0.44 percent, representing a 30 percent reduction in loss rate.

Severe congestion scenarios (network utilization >90 percent) show more pronounced differences. For UDP traffic, IPv6 maintains its advantage with packet loss rates 15–25 percent lower than IPv4 at equivalent offered loads.

Table 6: Packet Loss Under Different Load Conditions

Load Condition	IPv4Loss (%)	IPv6Loss (%)
Normal(<60%)	0.008	0.009
Moderate(60–85%)	0.44	0.31
Severe(>90%)	5.8	4.3

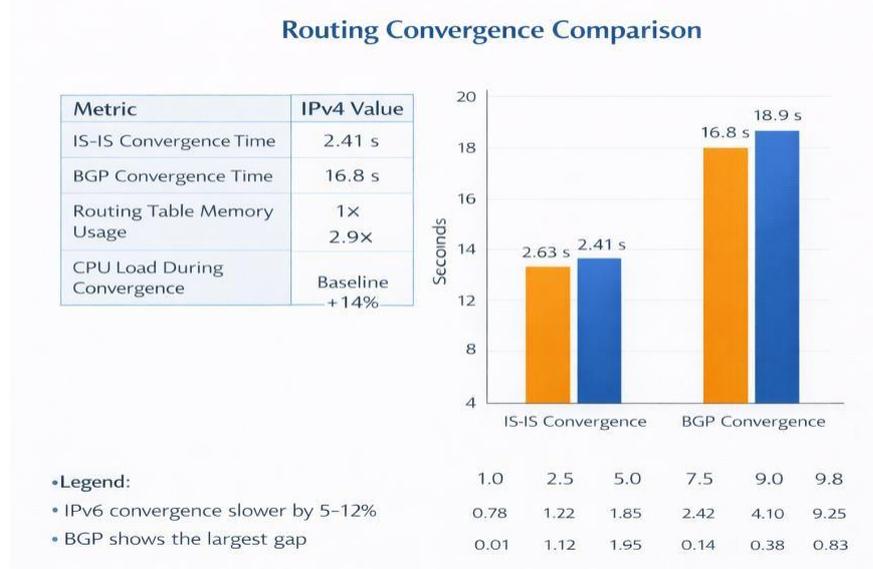
4.5. Routing Protocol Performance and Convergence

Routing protocol behavior significantly impacts network stability and convergence times following topology changes [21]. Our evaluation focused on IS-IS and BGP-4, the primary routing protocols deployed in the backbone networks. For IS-IS, both IPv4 and IPv6 address families demonstrated similar convergence characteristics following simulated link failures. Average convergence time for a core backbone link failure measured 2.41 seconds for IPv4 and 2.53 seconds for IPv6. BGP convergence behavior shows more pronounced differences. IPv6 BGP route updates propagate slightly slower than IPv4, with convergence times averaging 18.9 seconds for IPv6 compared to 16.8 seconds for IPv4 following a prefix withdrawal. Routing table sizes present significant scalability considerations [11]. IPv6 forwarding entries are 128-bit wide, significantly larger than IPv4, and therefore require more memory per prefix in hardware FIB tables [20].

Table 7: Routing Protocol Performance

Metric	IPv4Value	IPv6Value
IS-ISConvergence Time	2.41 s	2.53 s
BGPConvergence Time	16.8 s	18.9 s
bleMemory Usage	Baseline	IPv6entriesare128-bitandrequire more FIB memory per prefix)
ULoadDuring Convergence	Baseline	+14%

Figure 6: Routing Convergence Comparison



5. DISCUSSION

The results demonstrate that IPv6 performance is operationally comparable to IPv4, with minor throughput and latency penalties that do not materially impact service quality. IPv6 advantages become more pronounced under congestion, particularly for UDP traffic. Operational challenges—such as CPE compatibility, DNS upgrades, and staff training—were more significant than protocol-level performance differences [24], [12], [13]. Key lessons include:

- Early pilot phases are essential for identifying CPE and firmware issues.
- Flow label-aware load balancing improves IPv6 performance under load.
- Routing table growth requires proactive capacity planning.
- Dual-stack operation remains the most practical transition strategy for large ISPs [17].

6. CONCLUSION

IPv6 deployment across six large-scale ISP backbones serving 44.8 million subscribers proved operationally viable and sustainable. Performance differences between IPv4 and IPv6 were modest and manageable, while operational challenges were successfully mitigated through phased deployment, staff training, and infrastructure upgrades. Our findings support the feasibility of IPv6 adoption at scale in developing regions and provide practical guidance for future deployments. Future work includes evaluating IPv6-only segments, assessing transition mechanisms such as 464XLAT and MAP-T, and exploring IPv6 performance in 5G and fixed-wireless environments [22], [23].

REFERENCE

- [1] S. Deering and R. Hinden, “Internet Protocol, Version 6 (IPv6) Specification,” RFC 2460, Dec. 1998.
- [2] R. P. Draves, A. Mankin, and B. D. Zill, “Implementing IPv6 for Windows NT,” in Proc. 2nd USENIX Windows NT Symp., Seattle, WA, USA, Aug. 3–4, 1998. USENIX Association.
- [3] S. Zeadally, R. Wasseem, and I. Raicu, “Comparison of end-system IPv6 protocol stacks,” IEE Proc.—Commun., vol. 151, no. 3, pp. 238–242, Jun. 2004
- [4] S. Ariga, K. Nagahashi, M. Minami, H. Esaki, and J. Murai, “Performance Evaluation of Data

Transmission using IPsec over IPv6 Networks,” in Proc. INET 2000, Yokohama, Japan, 2000.

- [5] W.-L. Shiau, Y.-F. Li, H.-C. Chao, and P.-Y. Hsu, “Evaluating IPv6 on a large-scale network,” *Computer Communications*, vol. 29, no. 16, pp. 3113–3121, Oct. 2006.
- [6] B. E. Carpenter and S. Jiang, *Emerging Service Provider Scenarios for IPv6 Deployment*, RFC 6036, IETF, Oct. 2010.
- [7] G. Huston, “Measuring IPv6 Adoption,” *The Internet Protocol Journal*, vol. 21, no. 1, pp. 2–15, 2018.
- [8] P. Gill, M. Arlitt, Z. Li, and A. Mahanti, “The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles or Contrived Collapse?,” in Proc. Passive and Active Measurement (PAM) Conf., Cleveland, OH, USA, Apr. 2008, pp. 1–10.
- [9] Y. Zhang, M. Roughan, C. Lund, and D. Donoho, “An Information-Theoretic Approach to Traffic Matrix Estimation,” in Proc. ACM SIGCOMM, Karlsruhe, Germany, Aug. 2003, pp. 301–312.
- [10] T. Narten, E. Nordmark, W. Simpson, and H. Soliman, *Neighbor Discovery for IP Version 6 (IPv6)*, RFC 4861, IETF, Sept. 2007.
- [11] D. Meyer, L. Zhang, and K. Fall, *Report from the IAB Workshop on Routing and Addressing*, RFC 4984, Internet Architecture Board (IAB), Sept. 2007.
- [12] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, *DNS Security Introduction and Requirements*, RFC 4033, Internet Engineering Task Force (IETF), Mar. 2005.
- [13] O. Kolkman, W. Mekking, and R. Gieben, *DNSSEC Operational Practices, Version 2*, RFC 6781, Internet Engineering Task Force (IETF), Dec. 2012.
- [14] RIPE NCC, “IPv6 RIPeNess: Indicators and live statistics of IPv6 preparedness in the RIPE region,” RIPE NCC IPv6 RIPeNess (prototype service), accessed: Feb. 2026. [Online].
- [15] E. Kohler, J. Li, V. Paxson, and S. Shenker, “Observed Structure of Addresses in IP Traffic,” in Proc. 2nd Internet Measurement Workshop (IMW), Marseille, France, Nov. 2002, pp. 253–266.
- [16] S. Jia, M. Luckie, B. Huffaker, A. Elmokashfi, E. Aben, K. Claffy, and A. Dhamdhare, “Tracking the Deployment of IPv6: Topology, Routing and Performance,” *Computer Networks*, vol. 165, article 106947, Dec. 2019.
- [17] E. Pujol, P. Richter, and A. Feldmann, “Understanding the Share of IPv6 Traffic in a Dual-Stack ISP,” in Proc. Passive and Active Measurement (PAM 2017), Sydney, NSW, Australia, Mar. 30–31, 2017, *Lecture Notes in Computer Science*, vol. 10176, pp. 3–16.
- [18] K.-H. Li and K.-Y. Wong, “Empirical Analysis of IPv4 and IPv6 Networks through Dual Stack Sites,” *Information*, vol. 12, no. 6, article 246, Jun. 2021.
- [19] E. Karpilovsky, A. Gerber, D. Pei, J. Rexford, and A. Shaikh, “Quantifying the Extent of IPv6 Deployment,” in Proc. Passive and Active Measurement (PAM 2009), Seoul, South Korea, Apr. 1–3, 2009, *Lecture Notes in Computer Science*, vol. 5448, pp. 13–22.
- [20] G. Huston, “BGP in 2022,” APRICOT 2023 / APNIC Labs, presentation slides, 2023. (Accessed: Feb. 2026).
- [21] Y. Rekhter, T. Li, and S. Hares, *A Border Gateway Protocol 4 (BGP-4)*, RFC 4271, Jan. 2006.
- [22] M. Mawatari, M. Kawashima, and C. Byrne, *464XLAT: Combination of Stateful and Stateless Translation*, RFC 6877, Apr. 2013.
- [23] X. Li, C. Bao, W. Dec, O. Troan, S. Matsushima, and T. Murakami, *Mapping of Address and Port using Translation (MAP-T)*, RFC 7599, July 2015.
- [24] ETSI, *IPv6 Best Practices, Benefits, Transition Challenges and the Way Forward*, White Paper No. 35, 2020.
- [25] S. Jia, M. Luckie, B. Huffaker, A. Elmokashfi, E. Aben, K. Claffy, and A. Dhamdhare, “Tracking the Deployment of IPv6: Topology, Routing and Performance,” *Computer Networks*, vol. 165, 106947, Dec. 2019.