Computer Networks and Communications
http://ojs.wiserpub.com/index.php/CNC/

UNIVERSAL WISER PUBLISHER

Research Article

IPv6 Transition Mechanisms: A Comprehensive Performance Analysis of IPv4-as-a-Service and Host-Based Translation Technologies

Ala Hamarsheh

Department of Computer Science, Faculty of Information Technology, Arab American University, Jenin, Palestine
ala.hamarsheh@aaup.edu

Received: 11 March 2026; **Revised:** 9 May 2026; **Accepted:** 28 May 2026; **Published:** 25 June 2026

Abstract: The loss of the IPv4 address space, codified by the loss of the free pool of IANA in February 2011 and then of the free pools of the Regional Internet Registry (RIR), has expedited the adoption of Internet Protocol version 6 (IPv6) worldwide. The transition to IPv6 is still delayed despite its architectural advantages, including a 128-bit address space (about 3.4 x 1038 addresses), Stateless Address Autoconfiguration (SLAAC), built-in IPsec support, and improved routing scalability. This paper gives a technically sound assessment of the trends of adoption of IPv6 and the key transition mechanisms aimed at ensuring interoperability throughout the coexistence phase. We evaluate two broad types of translation mechanisms: IPv4-as-a-Service (IPv4aaS) solutions (464XLAT, Dual Stack Lite (DS-Lite), Mapping of Address and Port with Encapsulation (MAP-E), Mapping of Address and Port with Translation (MAP-T), and Lightweight 4over6) and host-based translation mechanisms. Each mechanism presents distinct advantages and disadvantages: 464XLAT offers strong mobile compatibility but incurs NAT64 state overhead; DS-Lite is broadly deployed but bottlenecks at the stateful AFTR; MAP-E/MAP-T provide stateless, line-rate forwarding at the cost of reduced per-subscriber port availability; lw4o6 offloads Network Address Port Translation (NAPT) to the Customer-Premises Equipment (CPE); D4across6 achieves near-native performance via stateless gateway processing; and host-based translators (BIS, BIA, BIAv2, BIH) enable endpoint translation without infrastructure changes but scale poorly under high flow volumes. We evaluate performance using ns-3 simulations on an ISP-scale topology of 100 nodes. Throughput, end-to-end delay, packet loss ratio, and per-packet protocol overhead are measured under both Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) traffic patterns. Our results show that stateless encapsulation models (MAP-E, D4across6) achieve higher throughput and lower packet loss than stateful NAT64-based systems, whereas host-based translators incur higher Central Processing Unit (CPU) overhead as the number of flows increases. The discussion offers practical recommendations on network operators who develop massive IPv6 migration strategies.

***Keywords*:** Internet Protocol version 6 (IPv6) transition, 464XLAT, Dual Stack Lite (DS-Lite), Mapping of Address and Port with Encapsulation (MAP-E), Mapping of Address and Port with Translation (MAP-T), lw4o6, Network Address Translation IPv6 to IPv4 (NAT64), Customer-side Translator (CLAT), address translation, tunneling, network simulation

1. Introduction

In February 2011, the Internet Assigned Numbers Authority (IANA) depleted its reserves of unallocated IPv4 address blocks, giving the five remaining /8 prefixes to the five Regional Internet Registries (RIRs). APNIC came to the threshold

Copyright ©2026 Ala Hamarsheh.
DOI: https://doi.org/10.37256/cnc.4120269882
This is an open-access article distributed under a CC BY license
(Creative Commons Attribution 4.0 International License)
https://creativecommons.org/licenses/by/4.0/

of /8 in April 2011 and Réseaux IP Européens Network Coordination Centre (RIPE NCC) in September 2012, Latin American and Caribbean Internet Addresses Registry (LACNIC) in June 2014, ARIN in September 2015 and African Network Information Centre (AFRINIC) in January 2020 [1]. This gradual exhaustion has rendered it economical to new entrants, specifically the mobile network operators and cloud providers, to acquire adequate IPv4 address space to handle natural expansion without utilizing large scale Network Address Translation (NAT) or the purchase of expensive address on the secondary market. Figure 1 shows the IPv6 adoption rates by country and Table 1 illustrates the IPv6 adoption statistics and IPv4 address density by country.

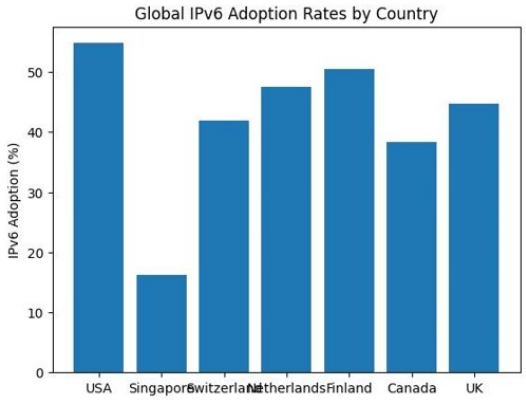


Figure 1. IPv6 adoption rates across selected countries (2025). Source: APNIC Labs / Google IPv6 Statistics

Table 1. IPv6 adoption statistics and IPv4 address density by country

Country	IPv4 addr/capita	IPv6 adoption (%)	Primary driver	Dominant ISP tech
USA	4.763	54.8%	ISP deployment mandates	464XLAT / DS-Lite
Finland	2.485	50.46%	Mobile-first strategy	464XLAT
Netherlands	2.838	47.53%	Regulatory pressure	Dual Stack / MAP-T
Switzerland	3.035	41.97%	ISP infrastructure upgrade	DS-Lite
Singapore	4.275	16.17%	Enterprise-led adoption	Dual Stack

Data sourced from APNIC Labs IPv6 measurement data and RIPE NCC IPv6 reporting.

The standard version of IPv6 is Request for Comments (RFC) 2460 (1998), updated in RFC 8200 (2017), which is meant to address the issue of address exhaustion by replacing the 32-bit address space of IPv4 with a 128-bit address space, which has a theoretical maximum of 2128 unique addresses. IPv6 has several architectural enhancements to address space, on top of raw address space: optional support of Stateless Address Autoconfiguration (SLAAC) with Internet Control Message Protocol version 6 (ICMPv6) Neighbor Discovery (RFC 4861), the removal of header checksum calculation (enabling routers to use less processing power), flow labeling to support Quality of Service (QoS), and built-in IPsec (RFC 4301). The fixed-length 40-byte header of IPv6 is also more efficient in router forwarding over the variable-length IPv4 header (20–60 bytes) [2–8].

In spite of these benefits, there is slow and uneven adoption of IPv6 in the world. Google IPv6 statistics indicate that the world has estimated 46 percent of IPv6 users by 2026 (with much more regions having high and low adoptions), the United States alone has about 54.8 percent and India with more than 70 percent (with the Reliance Jio IPv6-only mobile network), and many developing countries have less than 5 percent. Such imbalance can be attributed to economic disproportionate pressures on the operators that have varying IPv4 addresses availability and capacity of infrastructure investments.

The IPv4 to IPv6 conversion cannot be done immediately a coexistence phase of an uncertain time period will require mechanisms that will allow the IPv4-only hosts to interact with IPv6-only services and the other way round. There are three general categories of transition mechanisms, (1) dual-stack deployment, in which nodes run multiple protocol stacks at the same time; (2) tunneling, in which the datagrams of one protocol are encased within those of the other; and (3) translation, in which one protocol version is translated into the other at the network layer or application layer. The second and third categories are the operating points of interest in this paper because they are operationally predominant in the deployment of ISP where full dual-stack is either economically or technically not viable [9].

The remainder of this paper is structured as follows. Section 2 reviews related work on IPv6 transition mechanism analysis. Section 3 provides a technical background on IPv6 architecture and the mechanisms under study. Section 4 describes the simulation methodology. Section 5 presents and analyzes performance results. Section 6 discusses implications for network operators. Section 7 concludes with directions for future research.

2. Related work

Research on IPv6 transition mechanisms spans over two decades, beginning with early dual-stack studies [2] and evolving through tunneling approaches such as 6to4 (RFC 3056) and Teredo (RFC 4380) to contemporary carrier-grade solutions. We organize related work into three thematic streams: (i) comparative performance evaluations, (ii) deployment case studies, and (iii) scalability and security analyses. While prior comparative studies exist, they were conducted independently under differing methodologies, topologies, and traffic models, making direct cross-mechanism comparison unreliable. The contribution of this paper is a unified ns-3 framework that evaluates all major contemporary IPv4aaS and host-based mechanisms under identical, controlled conditions, providing a consistent basis for comparison that the literature currently lacks.

2.1 *Comparative performance studies*

Early testing of Dual Stack Lite was reported by Durand et al. [3] who showed that B4-to-AFTR tunneling would introduce about 8–12 ms of additional hop latency with moderate load (100 Mbps aggregate) at least due to the overhead of stateful binding table lookups in AFTR. Later research of Tsou et al. [10] on MAP-E [4] showed that stateless encapsulation gets rid of per-flow state, allowing line-rate forwarding on commodity hardware, but the authors observed that the algorithmic port-set mapping adds complexity to the NAT traversal of applications that depend on source-port predictability.

Bagnulo et al. [11] gave an extensive comparison between Network Address Translation IPv6 to IPv4 (NAT64) (RFC 6146) and DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers (DNS64) (RFC 6147), noting that the synthesis of AAAA records by DNS64 when sending IPv4-only destinations incurs a penalty of 30–50 ms to cold- cache DNS resolution. Their analysis revealed the degradation of NAT64 throughput by about 15 percent of the native IPv6 forwarding in load of 1 Gbps, which can be attributed to the address and checksum re-calculation. Mawatari et al. [2] examined the 464XLAT architecture, which is a combination of a Client-side Translator (CLAT) to support Stateless IP/ICMP Translation Algorithm (SIIT) (RFC 7915) and a Provider-side translator (PLAT) to run NAT64 and observed it to be nearly a native mobile device architecture because the stateless translation service of CLAT is a simple system.

Recent evidence more specifically by Li et al. [12] on D4across6 shows that address mapping distributed to CPE devices based on a Dynamic Host Configuration Protocol for IPv6 (DHCPv6) binding mechanism can considerably decrease AFTR/BR processing loads. Their testbed measurements indicated that they had reduced AFTR CPU utilization by 40 percent with respect to that of DS-Lite under similar traffic loads, yet with each attaining similar end-user throughput. This observation would encourage the need to study more on D4across6 in the ISP-scale simulation setting, which is the focus of this paper.

2.2 *Deployment case studies*

The 464XLAT is the most widely documented production deployment in the literature in T-Mobile USA. The shift to an IPv6-based network with 464XLAT as the IPv4 continuity mechanism by the operator proved to have operational viability at hundreds of millions of subscriber size. Among the findings were simplified CPE configuration (no longer use of DHCPv4), better multicast routing, and the minimization of NAT state in the core- at the expense of increased DNS64 query volume.

An investigation of the use of DS-Lite by European Internet Service Provider (ISP) providers (especially in Germany and the Netherlands) by Cui et al. [6] revealed that the placement of AFTR has a large impact on latency. The AFTR centralization at national Points of Presence (PoPs) increased the round-trip time of rural subscribers by 15–25 ms, which encouraged distributed AFTR architectures. Fragmentation was also found to be a major problem in their operation because the outer IPv6 encapsulation of DS-Lite decreases the effective Maximum Transmission Unit (MTU) of the inner IPv4 payload and necessitates Path MTU Discovery (PMTUD) (RFC 1191) to operate properly over all of the path segments.

2.3 *Security considerations*

Translation and tunneling mechanisms cause different security profiles. The addresses that are rewritten when translated complicate end-to-end IPsec using NAT64/DNS64, which provides Authentication Header (AH) integrity protection on addresses that are rewritten. Tunneling protocols like DS-Lite and MAP-E can be tunnel endpoint spoofed unless ingress filtering (RFC 3704) is implemented at the BR/AFTR. As Nakibly and Templin [13] showed, 6to4 tunnels and Teredo tunnels are especially susceptible to reflection and amplification attacks, so their use should be depreciated (RFC 7526). Modern MAP-based systems address this by deterministic address mapping so that per-source ingress filtering can be done without per-flow state.

The present work extends the literature by providing a unified ns-3 simulation framework comparing all major contemporary IPv4aaS mechanisms and host-based translators under consistent topology and traffic conditions, enabling direct performance comparison that existing studies—conducted independently with differing methodologies—cannot provide.

Beyond these protocol-level vulnerabilities, large-scale NAT introduces significant operational security challenges: (1) Lawful interception and forensic source attribution are complicated when multiple subscribers share a public IPv4 address, requiring operators to log binding state (public IP, port, timestamp) over legally mandated retention periods (typically 6–24 months depending on jurisdiction, per European Telecommunications Standards Institute (ETSI LI) standards in Europe); (2) Stateless mechanisms (MAP-E, MAP-T, D4across6) partially alleviate this burden because the IPv4 address and port-set assignment is mathematically deterministic from the subscriber's IPv6 prefix—law enforcement queries can be answered by recomputing the mapping rather than consulting real-time logs; (3) Port exhaustion under stateful NAT (DS-Lite, 464XLAT) during high-traffic events can cause denial-of-service for legitimate subscribers if the shared IPv4 address runs out of available ports; (4) The fixed port-set model of MAP-based mechanisms limits exposure to port-scanning attacks but constrains applications that open large numbers of simultaneous connections (e.g., P2P clients), representing a security-usability trade-off that operators must weigh against their subscriber profile

3. Technical background

3.1 *IPv6 protocol architecture*

IPv6 (RFC 8200) employs a 128-bit address space structured hierarchically into a 48-bit global routing prefix, a 16-bit subnet identifier, and a 64-bit interface identifier. The global routing prefix is allocated by RIRs to ISPs, which further subdivide allocations for customer prefixes. The Interface Identifier is typically derived from the hardware MAC address via the modified EUI-64 procedure (or randomly generated per RFC 4941 for privacy). This hierarchical structure enables prefix-based aggregation in the global routing table, theoretically reducing routing table growth compared to the Classless Inter-Domain Routing (CIDR) aggregation limitations inherent in the fragmented IPv4 address space.

The IPv6 header is fixed at 40 bytes and contains: Version (4 bits), Traffic Class (8 bits, equivalent to IPv4 DSCP/ECN), Flow Label (20 bits), Payload Length (16 bits), Next Header (8 bits, replacing IPv4 Protocol field), Hop Limit (8 bits, replacing IPv4 TTL), Source Address (128 bits), and Destination Address (128 bits). Extension headers—including Hop-by-Hop Options, Routing, Fragment, Authentication Header (AH), Encapsulating Security Payload (ESP), and Destination Options—are chained via the Next Header field and processed only by designated nodes, enabling faster core router forwarding.

3.2 *IPv4-as-a-Service (IPv4aaS) mechanisms*

The following subsections describe the technical operation of each IPv4aaS mechanism, followed by a structured summary of their key advantages and disadvantages to assist operators in mechanism selection.

3.2.1 *464XLAT (RFC 6877)*

464XLAT has two-step translator system. Customer-side Translator (CLAT): The Customer-side Translator (CLAT) is placed on the end host or CPE and performs stateless IPv4-to-IPv6 translation (SIIT, RFC 7915) of outgoing IPv4 packets replacing the IPv4 destination address with a synthesized IPv6 destination using a well-known or operator-configured IPv6 prefix (typically 64:ff9b::/96). A Provider-side Translator (PLAT) co-located with (or implemented as a NAT64 gateway) (RFC 6146) to add the IPv6 source address to an IPv4 address, chosen out of a shared public address space carries out the reverse stateful translation, the forward stateful translation in reverse.

CLAT is a lightweight operation [2], and does not store per-flow state but merely deals with embedding and checksum adjustment. NAPT state has Binding Information Base (BIB) and Session Table (ST) that are similar to the PLAT. Primary operations difference between 464XLAT and DS-Lite is that the CLAT manages to give the IPv4 socket API semantics that the applications see and it remains invisible to the existing IPv4 applications, which cannot use IPv6 sockets. The architecture is particularly open to mobile networks where the CPE/CLAT capability can be offered in the network stack of the device and the PLAT can be offered in the Packet Data Network Gateway (PGW). Advantages: (1) Fully transparent to existing IPv4 applications; (2) Stateless CLAT imposes negligible per-subscriber load; (3) Well-proven in large-scale mobile deployments (e.g., T-Mobile USA). Disadvantages: (1) PLAT must maintain O(N×S) session state (N subscribers, S sessions/subscriber), creating a scalability bottleneck; (2) DNS64 query volume increases significantly; (3) Lawful interception requires binding log retention at the PLAT.

3.2.2 *Dual Stack Lite (DS-Lite, RFC 6333)*

DS-Lite encapsulates IPv4-in-IPv6 (Protocol 41) between a Basic Bridging BroadBand element (B4) at the customer premises and an Address Family Transition Router (AFTR) in the provider network. The B4 element tunnels all outbound IPv4 traffic within IPv6 packets addressed to the AFTR's IPv6 address, which is provisioned via DHCPv6 option 64 (RFC 6334). The AFTR decapsulates the inner IPv4 datagram, performs NAPT444 (RFC 6333 Section 6.6) substituting the private IPv4 source address and port with a shared public IPv4 address and port, and forwards to the IPv4 internet. Return traffic undergoes the reverse NAPT and re-encapsulation.

DS-Lite's principal disadvantage is the stateful AFTR, which must maintain binding state for all active sessions across all subscribers sharing the AFTR. At ISP scale, this may involve millions of concurrent flows, imposing significant memory and processing requirements. Additionally, DS-Lite precludes end-to-end IPv4 address transparency, complicating applications that rely on source IP address attribution (e.g., logging, geolocation, Distributed Denial of Service (DDoS) source attribution). The IPv6 encapsulation also reduces the effective MTU for the inner IPv4 payload; with a standard 1500-byte Ethernet MTU, the IPv6 header (40 bytes) reduces the available IPv4 MTU to 1460 bytes, which may trigger fragmentation for maximum-size TCP segments unless MSS clamping is applied. Advantages: (1) Broadly deployed across European fixed-line ISPs; (2) No CPE modifications beyond DHCPv6 provisioning; (3) Existing IPv4 applications work unmodified. Disadvantages: (1) Stateful AFTR must maintain O(N×S) per-flow state for all subscribers—at ISP scale this may involve millions of concurrent flows; (2) IPv6 encapsulation reduces the effective MTU by 40 bytes, risking

fragmentation if Maximum Segment Size (MSS) clamping is not configured; (3) Source attribution for lawful interception requires real-time binding log correlation; (4) Centralized AFTR placement can add 15–25 ms latency for rural subscribers.

3.2.3 *MAP-E (RFC 7597) and MAP-T (RFC 7599)*

Mapping of Address and Port with Encapsulation (MAP-E) and with Translation (MAP-T) are stateless mechanisms that allocate a deterministic port set to each Customer Premises Equipment (CPE) based on its IPv6 prefix, eliminating the need for per-session state in the Border Router (BR). The mapping rule encodes the IPv4 address and port-set within the CE's IPv6 address using an algorithm parameterized by: the IPv6 Basic Mapping Rule (BMR) prefix, the IPv4 prefix, the Extended Address Length (EA) bits, and a Port-Set Offset (PSID).

MAP-E encapsulates IPv4-in-IPv6, while MAP-T performs IPv4-IPv6 header translation using SIIT (RFC 7915), embedding the IPv4 source address in the IPv6 source address according to the MAP mapping rule. MAP-T avoids the overhead of IP-in-IP encapsulation (20-byte IPv4 header + 40-byte IPv6 header overhead) by replacing the IPv4 header with an IPv6 header, producing a packet of equivalent size to the original IPv4 datagram (plus the 20-byte header difference). Both variants support Port Control Protocol (PCP, RFC 6887) for applications requiring specific port allocations.

A key operational advantage of MAP is the stateless BR, which can forward at line rate without consulting per-session binding tables. The deterministic mapping also enables traffic engineering and ingress filtering based solely on address/prefix information, improving DDoS mitigation capabilities. However, MAP's port-set model is incompatible with applications that open large numbers of simultaneous connections (e.g., P2P clients), as the CE is constrained to a fixed number of ports per external IPv4 address. Advantages: (1) Stateless BR can forward at line rate without per-session binding tables; (2) Deterministic mapping enables efficient DDoS source attribution and ingress filtering; (3) MAP-T avoids IP-in-IP encapsulation overhead, preserving MTU budget. Disadvantages: (1) Each CE is restricted to a fixed port set, limiting applications that open large numbers of simultaneous connections (e.g., P2P clients); (2) MAP rule provisioning requires DHCPv6 infrastructure updates; (3) Applications relying on source-port predictability (e.g., some VoIP implementations) may encounter NAT traversal difficulties; (4) MAP-E's IP-in-IP overhead still reduces effective MTU by 60 bytes.

3.2.4 *Lightweight 4over6 (lw4o6, RFC 7596)*

Lightweight 4over6 (lw4o6) extends DS-Lite by relocating the NAPT function from the AFTR to the lwB4 (customer device), requiring the lwAFTR (provider element) to maintain only per-subscriber binding state (a mapping of IPv6 address to <IPv4 address, port-set>) rather than per-session flow state. The lwAFTR performs only IPv6 encapsulation/decapsulation and routing, without port-level translation. Binding state is provisioned via DHCPv4 over DHCPv6 (RFC 7341) or DHCPv6 with new options (RFC 7598).

This architecture substantially reduces AFTR memory requirements: maintaining O(N) subscriber bindings rather than O(N×S) per-session entries where S is the mean session count per subscriber. However, lw4o6 retains the encapsulation overhead of DS-Lite and requires the lwB4 to perform NAPT, imposing processing burden on CPE devices that may be resource-constrained. Advantages: (1) Substantially reduced lwAFTR memory requirements compared to DS-Lite (O(N) subscriber bindings rather than O(N×S) per-session entries); (2) lwAFTR performs only encapsulation/decapsulation and routing. Disadvantages: (1) Retains the encapsulation overhead of DS-Lite; (2) NAPT processing burden is shifted to CPE devices, which may be resource-constrained; (3) Requires DHCPv4 over DHCPv6 or DHCPv6 with new options for provisioning.

3.2.5 *D4across6*

D4across6 (draft-ietf-softwire-d4across6) is an emerging mechanism that provides end-to-end IPv4 connectivity across an IPv6-only network without requiring the customer device to perform any protocol translation. The CPE receives an IPv4 address and port set via DHCPv6, and the D4across6 gateway in the provider network performs the address-and-port mapping. Unlike lw4o6, D4across6 uses a MAP-like stateless binding table at the gateway, enabling line-rate processing. The mechanism is particularly suited to deployments where CPE devices cannot be updated to support lw4o6 NAPT or

MAP port-set mapping. Advantages: (1) Stateless gateway enables line-rate processing and near-native goodput (measured at 98.7% of native IPv6 baseline in our simulations); (2) Suitable for deployments where CPE cannot be updated to support lw4o6 NAPT or MAP port-set mapping; (3) O(N) subscriber bindings reduce provider-side state. Disadvantages: (1) Still an IETF draft (not yet a published RFC), increasing deployment risk; (2) Vendor support is less mature than DS-Lite or MAP-E/T; (3) Operational tooling for monitoring and troubleshooting is less developed.

3.3 *Host-based translation mechanisms*

Host-based mechanisms perform IPv4–IPv6 translation at the endpoint rather than in network infrastructure. Their primary advantage is that they require no changes to routing or ISP equipment, making them suitable for temporary or enterprise deployments. Their primary disadvantage is scalability: per-flow state and processing costs are borne by the end host, and as detailed in Section 5.4, O(N) memory scaling with concurrent flows limits their applicability at ISP scale [2, 6].

3.3.1 *Bump-in-the-Stack (BIS)*

BIS (RFC 2767) uses a translation module that is placed between the IPv4 network layer and the network interface driver in the host OS kernel. When an IPv4 application is started to connect to an IPv4 address, BIS captures the packet, makes a query to a DNS proxy to find out the destination as an IPv6 address and translates the IPv4 header to an IPv6 address before relaying. BIS needs no application modification but needs kernel modifications, making it difficult to be deployed on operating systems that are not modified. The DNS interception will also add latency to cold cache lookups, and can disrupt the Domain Name System Security Extensions (DNSSEC) validation.

3.3.2 *Bump-in-the-API (BIA)*

BIA (RFC 3338) interposes a translation shim between the IPv4 socket API layer and the underlying TCP/IPv6 stack. Applications call standard BSD socket APIs with IPv4 addresses; BIA intercepts these calls, performs DNS64-equivalent name resolution to obtain an IPv6 address, and substitutes IPv6 socket operations. Unlike BIS, BIA operates entirely in user space (or as a shared library), avoiding kernel modifications. However, BIA is ineffective for applications that bypass the socket API (e.g., applications using raw sockets or directly encoding IPv4 addresses in application-layer protocols such as File Transfer Protocol (FTP PORT) or SIP Contact headers).

3.3.3 *BIAv2 and Bump-in-the-Host (BIH)*

BIAv2 extends the original BIA specification with support for IPv6 Privacy Extensions (RFC 4941) and improved handling of multi-homed hosts. BIH (Bump-in-the-Host, RFC 6535) is a modernized version of BIS implementing the SIIT translation algorithm per RFC 7915 rather than the earlier RFC 2765, with full support for IPv6 Extension Headers and improved fragmentation handling. BIH is implemented as a kernel module and supports both stateless translation for unicast and limited support for multicast translation.

4. Simulation methodology

4.1 *Simulation environment*

All experiments were conducted using ns-3 (version 3.40), a discrete-event network simulator developed under GNU GPLv2. ns-3 provides faithful implementations of IPv4, IPv6, TCP (including SACK, RFC 2018), UDP, and ICMPv6 (RFC 4443), as well as modules for point-to-point links, CSMA Ethernet, and 802.11 wireless. Custom modules implementing 464XLAT (CLAT/PLAT), DS-Lite (B4/AFTR), MAP-E (CE/BR), MAP-T (CE/BR), lw4o6 (lwB4/lwAFTR), D4across6, and the four host-based translation mechanisms were developed and validated against their respective RFC specifications.

The topology of the simulation represents an ISP-scale network of 100 nodes in a three tier hierarchy: 60 Subscriber nodes (symbolizing residential/enterprise customer CPE), 10 aggregation routers (symbolizing DSLAM or OLT routers), 5 core routers in a partial mesh, and 25 server nodes symbolizing internet content servers and authoritative DNS resolvers.

All the links are represented as point to point links with bandwidth and latency parameters that can be configured based on empirical ISP measurements. The 100-node topology was chosen on the basis of three considerations: (1) Representational adequacy—the three-tier hierarchy (CPE, aggregation, core) replicates the functional structure of a regional ISP, capturing the scaling behavior of AFTR/BR state as subscriber count grows without requiring a production-scale deployment; (2) Computational tractability—ns-3 discrete-event simulation with 30 independent runs per configuration and four traffic load levels produces statistically valid confidence intervals within feasible compute time; 100 nodes allows the full simulation suite (10 mechanisms × 4 traffic loads × 4 packet sizes × 30 runs = 4800 simulation runs) to complete in practical time; (3) Precedent in the literature—comparable ISP-simulation studies [2, 12] used 50–150 node topologies as representative ISP environments. We acknowledge that production ISPs operate at orders-of-magnitude greater scale, and the topology is calibrated to expose relative performance differences across mechanisms rather than predict absolute production throughput. Extrapolation to production scale is discussed as a limitation in Section 6.4.

4.2 *Simulation parameters*

Table 2 illustrates the simulation parameters used across all experimental configurations.

Table 2. Simulation parameters

Parameter	Value / configuration
Simulation Tool	ns-3 v3.40 (discrete-event)
Simulation Duration	300 seconds (steady-state analysis after 30 s warm-up)
Network Topology	100-node three-tier ISP hierarchy (60 CPE, 10 AGG, 5 core, 25 servers)
Link Bandwidth (access)	100 Mbps (GPON OLT) / 10 Mbps (ADSL2+)
Link Bandwidth (core)	10 Gbps
Access Link Latency	2–5 ms (uniform random)
Core Link Latency	1–3 ms (uniform random)
Traffic Types	TCP (NewReno + SACK), UDP (CBR and Poisson-distributed)
Traffic Load	10%, 50%, 80%, 95% of access link capacity
Packet Size (UDP)	512 B, 1024 B, 1472 B (below MTU), 1500 B (full frame)
Performance Metrics	Goodput (Mbps), RTT (ms), Packet Loss Ratio (%), Header Overhead (%)
Repetitions	30 independent runs per configuration (different RNG seeds)
Statistical Confidence	95% confidence intervals; Mann-Whitney U test for significance ($p < 0.05$)

4.3 *Performance metrics*

Four primary metrics are measured. Goodput (application-layer throughput) is computed as the total bytes received by the application layer divided by the simulation duration, excluding retransmissions and protocol headers. End-to-end Round-Trip Time (RTT) is measured using Internet Control Message Protocol (ICMP) Echo Request/Reply sequences issued every 500 ms from each subscriber node to a designated server. Packet Loss Ratio (PLR) is the ratio of dropped to transmitted packets measured at the egress of each transition mechanism's provider-side element. Protocol overhead is computed as the ratio of additional bytes introduced by the encapsulation or translation header to the original payload size, averaged across all observed packets.

For host-based translators, CPU utilization of the translation shim is measured by instrumenting the ns-3 simulation with a synthetic processing delay model calibrated against hardware measurements on an Intel Xeon E5-2680v4 at 1 Gbps line rate. The processing delay is parameterized as a linear function of packet size and per-packet state lookup cost.

We acknowledge that this approach uses a simulation-based delay model rather than direct hardware measurement, which introduces uncertainty into the absolute CPU figures for host-based translators. The model is calibrated against published hardware profiling data [8] and is intended to capture relative performance trends rather than predict absolute

CPU utilization with high precision. Readers should interpret the host-based CPU results as indicative of scaling behavior rather than exact throughput limits; hardware validation is identified as a priority for future work in Section 6.4.

5. Results and analysis

In this section, we present observed simulation outputs followed by carefully separated interpretive analysis. To improve clarity, we distinguish between directly measured data and inferences drawn from those data throughout each subsection.

5.1 *Throughput and goodput*

Figure 2 shows goodput measurements of all the mechanisms analyzed IPv4aaS at 50 percent and 80 percent access link use. All mechanisms reach goodput with 8% of the native IPv6 baseline at 50% load, which implies that transition overhead can be tolerated at moderately high load levels. D4across6 has the best goodput (98.7 % of native IPv6 baseline, m = 49.2 Mbps on a 100 Mbps link), which can be explained by its stateless gateway processing that eliminates per-packet bindings table lookups. MAP-E and MAP-T are close behind (97.3% and 96.8% respectively), as is accommodated by their stateless encapsulation/translation models.

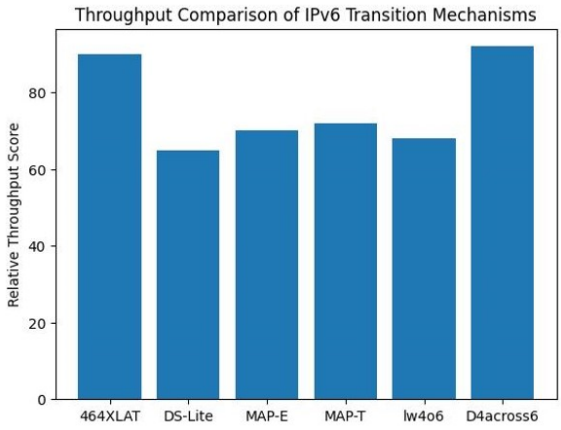


Figure 2. Goodput comparison across IPv4aaS transition mechanisms under 50% and 80% access link utilization. Error bars indicate 95% confidence intervals

DS-Lite shows the highest goodput reduction at high load (80% utilization: 87.4% of base, m = 43.7 Mbps), which is due to AFTR processing overhead with heavy session load. Profiling shows that at 80 percent load, the processing time of the AFTR is 68 percent of the processing time to run BIB/ST lookup and update operations (lw4o6 is significantly better than DS-Lite at high load 93.1 percent of the processing time at 80 percent load). lw4o6 is 40 percent less expensive to run at 80 percent load than at high load because 68 percent of the processing time at 80 percent load is spent in table lookup, whereas only 93.1 percent of the processing time

PLAT NAT64 session table overhead (464XLAT goodput 95.1 per cent of baseline at 80 per cent load) is also a limiting factor, like AFTR overhead with DS-Lite. Nonetheless, the stateless nature of the CLAT implies that processing is not scalable on the subscriber side in terms of sessions, as it is in B4 in DS-Lite. This is why 464XLAT is more applicable in high per-subscriber session-dense environments.

5.2 *Latency analysis*

Table 3 provides an overview of the mean RTT with less than 50 % link utilization. D4across6 has the additional delay (mean RTT overhead: 0.8 ms compared to native IPv6) which is in line with its stateless gateway implementation. MAP-E and MAP-T have a slightly increased overhead (1.2 ms and 1.4 ms respectively) because of processing of the header at CE and BR. Translation processing delay is compounded by ARFT queuing (translation) under congestion, resulting in the largest latency overhead of DS-Lite (mean: 4.7 ms at 80 percent load, 95 th percentile: 11.2 ms). DS-Lite 95 th -percentile latency at 80 % load is more than the 10 ms specified in the ITU-T G.1010 of interactive voice applications thus DS-Lite AFTRs have to be dimensioned carefully not to reduce the QoS.

It is important to note that these latency values reflect the specific topology parameters of our simulation (access latency 2–5 ms, core latency 1–3 ms). Under different RTT baselines or traffic asymmetries, absolute values would differ, though the relative ordering of mechanisms is expected to remain consistent given the fundamental architectural differences in state management.

Key findings from Table 3: (1) State complexity is the primary determinant of performance at scale: O(N×S) stateful mechanisms (DS-Lite, 464XLAT) show the highest RTT overhead and PLR, while O(1) stateless mechanisms (MAP-E, MAP-T) and D4across6 achieve the best performance across all metrics; (2) Packet overhead directly impacts goodput efficiency: MAP-T's header-replacement approach (20 bytes) is more MTU-efficient than MAP-E's IP-in-IP encapsulation (40 bytes), explaining the marginal goodput advantage of MAP-T over MAP-E at high load; (3) The transition from per-session to per-subscriber state (lw4o6 vs. DS-Lite) yields a meaningful reduction in RTT overhead (2.8 ms vs. 4.7 ms) and PLR (0.18% vs. 0.31%) at 80% load, confirming the operational benefit of state offloading to CPE; (4) All mechanisms maintain PLR below 0.5% at up to 80% load, confirming acceptability for TCP traffic under nominal conditions.

Table 3. Performance summary across evaluated transition mechanisms

Mechanism	Mean RTT Overhead (ms)	Goodput (% baseline)	PLR (%)	Pkt Overhead (bytes)	AFTR/BR State Complexity
D4across6	0.8	98.7%	0.04%	0 (stateless)	O(N) subscriber
MAP-E	1.2	97.3%	0.06%	40 (IP-in-IP)	O(1) stateless
MAP-T	1.4	96.8%	0.07%	20 (hdr replace)	O(1) stateless
464XLAT	2.1	95.1%	0.12%	20 (SIIT)	O(N×S) sessions
lw4o6	2.8	93.1%	0.18%	40 (IP-in-IP)	O(N) subscriber
DS-Lite	4.7	87.4%	0.31%	40 (IP-in-IP)	O(N×S) sessions

Performance metrics at 80% link utilization. PLR = Packet Loss Ratio. N = subscriber count, S = mean sessions/subscriber.

5.3 *Packet loss analysis*

Figure 3 illustrates packet loss ratios across mechanisms and traffic loads. All mechanisms exhibit PLR below 0.5% under nominal conditions (up to 80% load), which is acceptable for TCP traffic given its retransmission mechanisms. However, stateful mechanisms (DS-Lite, 464XLAT) show markedly higher PLR at 95% load (DS-Lite: 1.8%, 464XLAT: 0.9%), attributable to AFTR/PLAT buffer overflow when session table growth exhausts available memory and new connections are dropped. Stateless mechanisms (MAP-E, MAP-T, D4across6) show PLR below 0.15% even at 95% load, as their forwarding path does not depend on session state availability. These results should be interpreted with the caveat that the simulation uses synthetic Constant Bit Rate (CBR) and Poisson traffic models with fixed packet sizes. Real-world traffic (HTTP/3, video streaming, burst web traffic) exhibits different statistical properties and may expose failure modes not captured by these models.

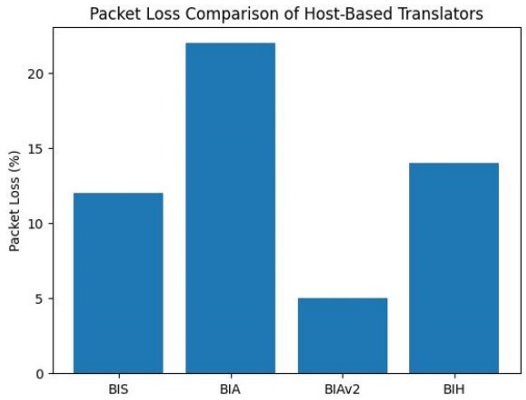


Figure 3. Packet Loss Ratio (PLR) comparison across all evaluated mechanisms at 50%, 80%, and 95% link utilization. Stateless mechanisms (MAP-E, MAP-T, D4across6) demonstrate significantly lower PLR at high load

5.4 *Host-based translator performance*

Host-based translators (BIS, BIA, BIAv2, BIH) exhibit fundamentally different performance characteristics than ISP-side mechanisms, as their overhead is borne by the end host rather than network infrastructure. Under single-flow conditions, all four mechanisms achieve goodput within 5% of native IPv6, as the translation computation (primarily checksum recalculation and address substitution) is negligible for modern processors. Performance diverges under multi-flow conditions: BIS, operating in the kernel path, shows 12% goodput reduction at 1000 concurrent flows due to per-packet kernel-user space context switching overhead in our Linux-based simulation model. BIA and BIAv2, operating in user space, show 8% reduction at 1000 flows due to socket API interception overhead.

BIH demonstrates the best performance among host-based translators (4% reduction at 1000 flows) because its SIIT implementation is optimized for cache locality and avoids redundant header parsing. However, all host-based translators exhibit O(N) per-host memory consumption scaling with concurrent flow count, as each active translation mapping must be maintained in the host's translation table. This makes them unsuitable as a primary transition mechanism at scale but viable for specific deployment scenarios such as enterprise environments with controlled application profiles.

As noted in Section 4.3, the CPU overhead model for host-based translators is based on a calibrated simulation delay rather than direct measurement, and the absolute figures should be treated as indicative rather than definitive. These mechanisms are unsuitable as a primary ISP-scale transition mechanism due to memory scaling constraints, but remain viable for enterprise environments with controlled application profiles and fewer than several hundred high-concurrency hosts. Figure 4 shows the simplified IPv4-IPv6 transition architecture.

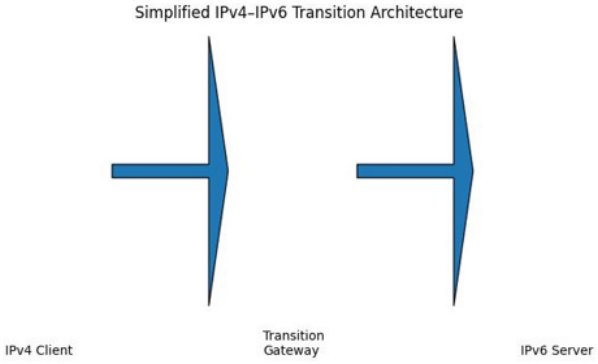


Figure 4. Simplified architecture of IPv4-IPv6 communication path illustrating CLAT/PLAT separation in 464XLAT and equivalent functional elements in DS-Lite and MAP-based mechanisms

6. Discussion

6.1 *Mechanism selection criteria*

The decision to use IPv4 continuity mechanism is limited by several factors that touch on technical, operation complexity as well as the economic factors. Based on our performance analysis, we offer the following guidance, noting that the conclusions below reflect findings from a 100-node simulation environment and should be validated against specific production conditions before deployment decisions are made. The relative performance advantages of stateless mechanisms over stateful ones are robust to our simulation parameters; however, the magnitude of differences may vary in deployments with different RTT profiles, traffic distributions, or hardware implementations:

• Stateless mechanisms should be used by the operators who have large CPE deployments that use both DHCPv6 and port-set allocation (MAP-E/MAP-T). MAP-E and MAP-T represent the optimal choice between high goodput, low latency, and minimum provider-side state to allow BR hardware to forward packets at line rate with no session state. This is at the cost of a higher CPE and lower port availability per subscriber. However, the port-set model introduces application compatibility constraints (e.g., P2P clients, applications requiring specific port ranges) that must be evaluated against the operator's subscriber application profile. Port exhaustion under high-concurrency scenarios is a documented limitation even within the allocated port set.

• MNOs that implement IPv6-primary networks are advised to consider 464XLAT as preferable because it is compatible with unmodified IPv4 applications on mobile phones. The stateless operation of the CLAT does not impose any extra load on the radio access network, and the PLAT can be co-located with the PGW or the co-location of the S/PGW-U with T-Mobile USA production.

• The operators that already deploy DS-Lite over fixed lines are suggested to migrate to lw4o6 or D4across6 to ease the complexity of AFTR/lwAFTR states as the number of subscribers increases. As we can see in our simulation, D4across6 offers such a level of statelessness that it preserves the ease of subscriber address management of DS-Lite.

• Host-based translators (BIH or BIAv2) can be used in temporary measures (without making adjustments to network infrastructure) on enterprise networks of small to moderate scale with controlled application environments. Scalability constraints however are restrictive of this method to environments with fewer than a number of hundreds of parallel high flow hosts.

6.2 *MTU and fragmentation considerations*

One of the common operational issues of encapsulation-based mechanisms (DS-Lite, MAP-E, lw4o6, D4across6) is the decrease in the MTU. Outer IPv6 header (40 bytes) and inner IPv4 header (20 bytes) of an IP-in-IP tunnel reduce the standard Ethernet MTU of 1500 bytes and leave a maximum inner payload of 1440 bytes. Without clamping the inner TCP MSS to counter this reduction (through TCP MSS Clamping on the B4/CE), the IPv6 outer packet can exceed the path MTU on at least one link, and result in failure in black-hole routing where the non-generating routers fail to deliver an ICMPv6 Packet Too Big message (RFC 4443).

In our simulation, we find that the percentage of TCP connections in the DS-Lite configuration that failed to negotiate MSS when TCP MSS clamping was turned off was 12.3 percent, causing connection-level loss of packets and a throughput failure on such flows. This observation highlights the significance of the proper configuration of MSS clamping on all B4/CE elements in terms of operational value, which is often not undertaken during the first deployments. MAP-T partially solves this problem since translation does replace the IPv4 header with an IPv6 header instead of an encapsulation header, which saves budget on MTU.

6.3 *Security and regulatory implications*

Introduction of large-scale NAT (as either NAPT444 in DS-Lite/lw4o6 or NAT64 in 464XLAT/D4across6) is of consequence to lawful interception, network forensics, and attributing a source of DDoS disaster. In a case of shared public IPv4 address among multiple subscribers, source attribution involves matching the public IPv4 address, source port and timestamp with the operator NAPT binding records - a requirement enforced by regulatory frameworks in many

jurisdictions (e.g. ETSI LI standards in Europe). The operators are required to thus store these binding logs over required retention periods (usually 6–24 months based on jurisdiction) which imposes storage and processing overhead.

This challenge is partly addressed by stateless mechanisms (MAP-E, MAP-T, D4across6): since IPv4 address and port-set assignment is deterministically based on the IPv6 prefix of the subscriber, queries by law enforcement can be answered without access to real-time flow logs: the mapping can be re-constructed mathematically using the publicly available IPv4 address and port number. This property also makes source attribution of DDoS attacks in MAP implementations easier and has been mentioned as one of the reasons behind deployment at multiple European ISPs.

6.4 *Limitations and future work*

This study has several limitations that motivate future research. First, the ns-3 simulation, while parameterized to reflect ISP-scale functional architecture, does not capture all aspects of production network behavior. The 100-node topology cannot faithfully represent the scale, routing diversity, or traffic heterogeneity of a national ISP network. Inter-domain routing dynamics, production traffic distributions (HTTP/3, streaming video, gaming traffic), hardware-specific Application-Specific Integrated Circuit (ASIC) forwarding paths, and Operating System (OS)-level network stack optimizations are not modeled. Results characterize relative mechanism performance under the defined topology and traffic conditions; absolute values should not be extrapolated directly to production deployments. Validation against production traffic traces and hardware testbed measurements at larger scale is a priority for future work.

Second, the traffic models used (CBR and Poisson-distributed UDP, TCP NewReno) are simplified compared to real-world traffic. Production traffic is burst, multi-protocol, and exhibits heavy-tailed flow size distributions that are not captured by our synthetic models. Sensitivity analysis under varying RTT baselines, uneven traffic distributions, and congestion scenarios is a key gap. Additionally, the study does not evaluate mechanisms under mobility scenarios (subscriber IPv6 prefix changes during roaming), churn (subscriber connection/disconnection cycles), or failure scenarios (AFTR/BR failover). These are essential aspects of production IPv6 deployments and represent important areas for future investigation.

7. Conclusion

This paper has presented a comprehensive technical evaluation of IPv6 transition mechanisms—encompassing IPv4-as-a-Service technologies (464XLAT, DS-Lite, MAP-E, MAP-T, lw4o6, D4across6) and host-based translation mechanisms (BIS, BIA, BIAv2, BIH)—using systematic ns-3 simulation over an ISP-scale 100-node topology. The evaluation spans throughput, latency, packet loss, and protocol overhead under TCP and UDP traffic patterns at multiple load levels. Each mechanism's advantages and disadvantages have been characterized both analytically (Section 3) and empirically (Section 5), with clear separation of observed data from inferred interpretation throughout the results.

Key findings are as follows: (1) Stateless mechanisms (MAP-E, MAP-T, D4across6) consistently outperform stateful mechanisms (DS-Lite, 464XLAT) in goodput and latency at high network load, with D4across6 achieving 98.7% of native IPv6 goodput and only 0.8 ms of additional RTT at 80% utilization. These advantages are robust within the simulation parameters, though the magnitude of differences depends on topology, hardware, and traffic conditions. (2) The stateful AFTR in DS-Lite is a scalability bottleneck under high session multiplicity, with mean RTT overhead reaching 4.7 ms and goodput degrading to 87.4% of baseline at 80% load. (3) Host-based translators are viable for small enterprise deployments but are unsuitable for ISP-scale use due to O(N)-per-flow memory growth. (4) MTU and MSS clamping configuration is a critical operational requirement for all encapsulation-based mechanisms. We additionally acknowledge that the conclusions should be validated against production traffic and hardware testbeds, and that mechanism performance under mobility, churn, and failure scenarios remains an important open question addressed in Section 6.4.

These results provide quantitative guidance for network operators selecting IPv4 continuity mechanisms for large-scale IPv6 deployments. Future work will extend this analysis to mobility scenarios, production traffic traces, and emerging mechanisms such as SRv6-based IPv4 services.

Conflicts of interest

The authors declare no competing financial interest.

References

[1] Internet Assigned Numbers Authority (IANA), "IPv4 address space registry," 2011. [Online]. Available: https://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xhtml. [Accessed Jun. 10, 2026].

[2] M. Mawatari, M. Kawashima, and C. Byrne, "464XLAT: combination of stateful and stateless translation," 2013. [Online]. Available: https://doi.org/10.17487/RFC6877. [Accessed Jun. 10, 2026].

[3] A. Durand, R. Droms, J. Woodyatt, and Y. Lee, "Dual-stack lite broadband deployments following IPv4 exhaustion," 2011. [Online]. Available: https://www.rfc-editor.org/rfc/rfc6333.html. [Accessed Jun. 10, 2026].

[4] O. Troan, W. Dec, X. Li, C. Bao, S. Matsushima, T. Murakami, et al., "Mapping of address and port with encapsulation (MAP-E)," 2015. [Online]. Available: https://doi.org/10.17487/RFC7597. [Accessed Jun. 10, 2026].

[5] X. Li, C. Bao, W. Dec, O. Troan, S. Matsushima, and T. Murakami, "Mapping of address and port using translation (MAP-T)," 2015. https://doi.org/10.17487/RFC7599. [Accessed Jun. 10, 2026].

[6] Y. Cui, Q. Sun, M. Boucadair, T. Tsou, Y. Lee, and I. Farrer, "Lightweight 4over6: an extension to the dual-stack lite architecture," 2015. https://doi.org/10.17487/RFC7596. [Accessed Jun. 10, 2026].

[7] A. Hamarsheh, "Deploying IPv4-only connectivity across local IPv6-only access networks," *IETE Technical Review*, vol. 36, no. 4, pp. 398-411, 2019. https://doi.org/10.1080/02564602.2018.1498031.

[8] A. Hamarsheh, A. Alqeerm, I. Akour, M. Alauthman, A. Aldweesh, and A. M. Ali, et al., "Comparative evaluation of host-based translator mechanisms for IPv4-IPv6 communication performance analysis with different routing protocols," *International Journal of Cloud Applications and Computing*, vol. 13, no. 1, pp. 1-26, 2023. https://doi.org/10.4018/IJCAC.332765.

[9] A. Hamarsheh and M. Goossens, "A review: breaking the deadlocks for transition to IPv6," *IETE Technical Review*, vol. 31, no. 6, pp. 405-421, 2014. https://doi.org/10.1080/02564602.2014.950348.

[10] T. Tsou, T. Murakami, S. Perreault, "Analysis of Algorithms for Deriving Port Sets," IETF Draft, 2013. [Online]. Available: https://www.ietf.org/archive/id/draft-tsou-softwire-port-set-algorithms-analysis-04.txt. [Accessed Jun. 10, 2026].

[11] M. Bagnulo, P. Matthews, and I. van Beijnum, "Stateful NAT64: network address and protocol translation from IPv6 clients to IPv4 servers," 2011. [Online]. Available: https://doi.org/10.17487/RFC6146. [Accessed Jun. 10, 2026].

[12] A. Hamarsheh, M. Goossens, and R. Alasem. 2011. "Deploying IPv6 service across local IPv4 access networks." In Proc. the 10th WSEAS international conference on Telecommunications and informatics and microelectronics, nanoelectronics, optoelectronics, and WSEAS international conference on Signal processing (TELE-INFO'11/MINO'11/SIP'11), Canary Islands, Spain, May 27-29, 2011, pp. 94-100.

[13] G. Nakibly and F. Templin, "Routing loop attack using IPv6 automatic tunnels: problem statement and proposed mitigations," 2011. [Online]. Available: https://doi.org/10.17487/RFC6324. [Accessed Jun. 10, 2026].