Alqalam Journal of Medical and Applied Sciences. 2026;9(6):1823-1831
https://doi.org/10.54361/ajmas.269650

Original article

NAT-Based Address Conservation Mechanisms: Performance Transparency, and End-to-End Architectural Implications

Salah Abeid, Reyad Abulajras, Nuredin Ahmed

Department of Computer Engineering, Faculty of Engineering, University of Tripoli, Libya
Email. Nu.ahmed@uot.edu.ly

Abstract
The exhaustion of the IPv4 address space has driven the widespread deployment of Network Address Translation (NAT) and its derivatives as primary address conservation mechanisms. While these technologies have successfully prolonged the usability of the IPv4 Internet, they introduce fundamental trade-offs across three critical dimensions: performance, transparency, and architectural integrity. This paper provides a comprehensive survey and analysis of NAT-based address conservation mechanisms, including traditional NAT/NAPT, Carrier-Grade NAT (CGNAT/NAT444), Address plus Port (A+P), Dual-Stack Lite (DS-Lite), Mapping of Address and Port (MAP), and the 4+4 architecture. We classify these mechanisms according to the location of the address sharing function, state storage requirements, and traversal methods. We then systematically evaluate each mechanism's performance characteristics, transparency implications, and compliance with the Internet's end-to-end principle. Our analysis reveals that while stateful approaches offer immediate deployability at the cost of scalability and transparency, stateless and hybrid mechanisms present promising alternatives that better preserve architectural principles at the expense of increased complexity. The paper concludes with recommendations for future research directions and deployment strategies in the transition toward IPv6. Index Terms: Network Address Translation (NAT), Carrier-Grade NAT (CGNAT), address conservation, IPv4 address exhaustion, end-to-end principle, network transparency, performance evaluation, IPv6 transition.
Keywords. ICMP network security, Path MTU Discovery, rate limiting IPv6.

Introduction

The exponential growth of the Internet has precipitated an unprecedented challenge: the exhaustion of the IPv4 address space. First projected by the IETF Address Lifetime Expectations (ALE) Working Group in 1994, the depletion of available IPv4 addresses has fundamentally altered the architecture and operation of the global Internet. In response to this scarcity, Network Address Translation (NAT) emerged as an incremental, pragmatic solution that enables address reuse without requiring comprehensive upgrades to hosts and routers. NAT operates by modifying IP address and port information in packet headers as traffic traverses network boundaries, allowing multiple hosts within a private addressing realm to share a small pool of public IP addresses. The fundamental premise of RFC 1631 was to conserve IPv4 addresses by permitting only a subset of hosts within a domain to communicate directly with the external Internet. While NAT has been remarkably successful in mitigating address depletion, enabling continued operation even in regions with acute IP address shortages, its widespread adoption has come at a high cost [1].

The architectural implications of NAT extend far beyond address conservation. As documented in RFC 2993, NAT fundamentally affects the transparency of end-to-end connectivity for transports that rely on consistency of the IP header, as well as for protocols that carry address information in locations beyond the IP header. The end-to-end principle—the foundational architectural tenet that certain functions can only be performed at the endpoints, and that the network should serve as a simple datagram service—is directly challenged by the insertion of stateful address translation devices in the network path [2][8]. The evolution of NAT-based conservation mechanisms has progressed through several generations. Traditional NAT and Network Address and Port Translation (NAPT) were initially deployed at enterprise and residential boundaries. As IPv4 exhaustion intensified, Internet Service Providers (ISPs) began deploying Carrier-Grade NAT (CGNAT), also known as NAT444 or Large-Scale NAT (LSN), which introduces an additional layer of address translation within the provider's network. More recent proposals have explored stateless approaches such as Address plus Port (A+P), encapsulation-based mechanisms like MAP-E and MAP-T, and hybrid architectures including DS-Lite and the 4+4 architecture [3]. This paper makes the following contributions:

1. A comprehensive classification of NAT-based address conservation mechanisms along multiple dimensions
2. A systematic analysis of performance characteristics and trade-offs across mechanisms
3. An evaluation of transparency implications for end hosts and applications
4. An assessment of architectural impact on the end-to-end principle
5. A comparative framework to guide mechanism selection in different deployment contexts

The remainder of this paper is organized as follows. Section II provides background on the IPv4 address exhaustion problem and the fundamental operation of NAT. Section III presents a classification framework for addressing conservation mechanisms. Section IV analyzes performance characteristics. Section V examines transparency implications. Section VI discusses architectural implications for the end-to-end

Received*: 28-03-2026 -* ***Accepted****: 26-05-2026 -* ***Published****: 01-06-2026*

principle. Section VII presents a comparative analysis and trade-off discussion. Section VIII concludes with future research directions [4].

Methodology

The IPv4 Address Exhaustion Problem

The Internet's addressing architecture, defined by IPv4, provides approximately 4.3 billion unique addresses. This allocation, while seemingly vast, has proven insufficient to accommodate the explosive growth of Internet-connected devices. The situation was exacerbated by the initial class-based addressing structure, which resulted in significant address space fragmentation and inefficiency [5]. Despite the development of IPv6 as a long-term solution beginning in 1994, deployment has proceeded at a pace insufficient to outpace address depletion. The IANA central pool of IPv4 addresses was exhausted in 2011, followed by depletion at the Regional Internet Registry (RIR) level in subsequent years. This reality has forced network operators to seek immediate, deployable solutions that extend the utility of the existing IPv4 infrastructure [6].

Network Address Translation: Principles and Operation

Network Address Translation, as originally specified in RFC 1631, provided pure address translation: the algorithmic mapping of one IPv4 address to a different IPv4 address. This basic form allowed hosts on either side of a NAT to initiate connections, though it offered limited conservation benefits. The more significant development was Network Address and Port Translation (NAPT), which enables the multiplexing of multiple private addresses onto a single public address through the use of transport-layer port numbers. NAPT operates by creating and maintaining a mapping state between (private IP, private port) pairs and (public IP, public port) pairs. When a packet traverses the NAT from the private to the public realm, the source address and port are rewritten according to the mapping table; reverse translation occurs for returning packets. NAT devices implement varying criteria for creating, preserving, and removing mapping state, and employ different filtering and forwarding rules. Common behaviors include:

Endpoint-independent filtering

Allows inbound traffic from any external endpoint to a mapped port.

Address-dependent filtering

Restricts inbound traffic to the specific external address that received outbound traffic.

Address and port-dependent filtering

Further restricts inbound traffic to the exact external address and port pair.

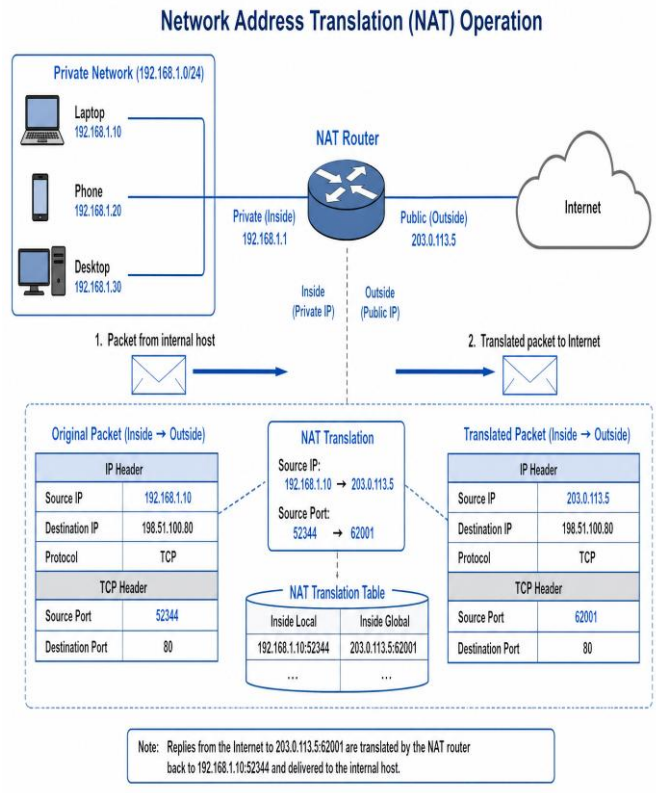

Figure 1. Fundamental NAT Operation

Received*: 28-03-2026 -* ***Accepted****: 26-05-2026 -* ***Published****: 01-06-2026*

The End-to-End Principle and Its Significance

The end-to-end principle, articulated in the foundational design of the Internet, holds that reliability and intelligence should reside at the end nodes of the network rather than within the network infrastructure itself. Under this model, the network provides a simple, transparent datagram service that moves bits between endpoints, each identified by globally unique and routable addresses [7,8].
This design has profound implications. It enables:

Application innovation

New protocols and applications can be deployed without network modifications.

Competitive neutrality

Network operators compete on bandwidth and reliability rather than on application-specific capabilities.

Security

End-to-end encryption and authentication can be implemented without network intermediaries.

Resilience

The network remains simple and robust while intelligence at the edges adapts to failures [9]. NAT fundamentally disrupts this model by modifying packet headers, introducing state in the network, and creating address realms that are not globally unique.

Results And Discussion

The landscape of NAT-based address conservation mechanisms encompasses a diverse set of approaches. Following the classification framework proposed in the literature, we categorize mechanisms along multiple dimensions [10].

Location of the Address Sharing Function

The address sharing function can be situated in different network locations, each with distinct implications:

Customer Premises Equipment (CPE)

In this model, address translation occurs at the customer's edge device. This is the traditional NAT deployment pattern, where the home router or enterprise gateway performs NAPT. The customer retains control over translation policies, including port forwarding and mapping behavior. A+P mechanisms exemplify this approach, where the CPE implements port-restricted NAPT or directly hosts A+P-capable endpoints.

Service Provider Gateway (CGN)

Carrier-Grade NAT places the address sharing function within the ISP's network infrastructure. This centralizes translation for large numbers of customers, enabling more efficient address utilization at the cost of introducing a critical network function under the provider's control. CGNAT and NAT444 operate in this mode.

Hybrid (CPE and Gateway)

Some mechanisms distribute address-sharing responsibilities between the CPE and the gateway. DS-Lite, for instance, places a lightweight tunnel endpoint (B4 element) at the CPE and a centralized tunnel concentrator (AFTR) in the provider network.

Dimension Significance

The location of the sharing function critically influences end-to-end transparency. A+P mechanisms, by placing translation at the customer premises, better preserve the Internet's end-to-end principle. Conversely, CGN deployments introduce an opaque translation layer between the customer and the public Internet, exacerbating transparency issues.

State Storage Requirements

The volume and granularity of state maintained by translation devices represent a fundamental performance and scalability determinant:

Per-Flow Stateful

Traditional NAPT maintains state on a per-flow basis, with each TCP, UDP, or ICMP session requiring a separate entry in the translation table. This approach provides maximum flexibility but imposes significant memory and processing requirements, particularly under high concurrent session loads. The state table experiences high churn as flows are short-lived, requiring efficient aging and garbage collection mechanisms.

Received*: 28-03-2026 -* ***Accepted****: 26-05-2026 -* ***Published****: 01-06-2026*

Per-Allocation Stateful
Some mechanisms allocate port ranges or sets to customers rather than managing individual flows. This approach reduces state granularity and associated overhead. The gateway stores mappings from IPv4 addresses and port-sets to tunnel identifiers, IPv6 prefixes, or CPE addresses. This represents a middle ground between fully stateful and stateless operation.

Stateless
Stateless mechanisms require no per-flow or per-allocation state in the gateway. A+P, for example, achieves statelessness by using port range bits from the TCP/UDP header as additional endpoint identifiers. The A+P gateway does not need to track every flow, as the port range deterministically identifies the customer. MAP-E and MAP-T similarly provide stateless address and port reuse.

Dimension Significance
State storage requirements directly impact scalability, cost, and performance. Stateless mechanisms offer superior scalability and eliminate state synchronization challenges in distributed deployments, at the cost of increased complexity in address and port allocation.

Traversal Method Through the Access Network
The method by which packets traverse between the CPE and gateway encompasses several approaches:

Routing
The simplest traversal method, routing, involves no packet header manipulation. IPv4 and IPv6 packets are carried natively from source to destination. This approach is characteristic of traditional NAT deployments where the public Internet routes packets natively.

Tunneling
Tunneling encapsulates packets within additional IP headers, allowing original packets to traverse non-native networks intact. DS-Lite encapsulates IPv4 packets in IPv6 tunnels between the CPE (B4 element) and the tunnel concentrator (AFTR). MAP-E similarly employs encapsulation.

Double Address Family Translation
This approach leverages stateless IP/ICMP translation (stateless NAT64) to translate between IPv4 and IPv6 header formats. MAP-T exemplifies this method, performing stateless dual translation.

Reversible Header Translation
Some mechanisms employ header translation that can be reversed at the egress, preserving end-to-end transparency while enabling address sharing.

Summary of Major Mechanisms
(Table 1) summarizes the major NAT-based address conservation mechanisms according to the classification dimensions described above.

Table 1. Classification of NAT-Based Address Conservation Mechanisms

Mechanism	Location	State Storage	Traversal Method	Key Characteristic
Traditional NAT/NAPT	CPE	Per-flow	Routing	Widely deployed, simple
CGNAT/NAT444	Gateway	Per-flow	Routing	ISP-scale, double NAT
A+P	CPE	Stateless	Routing	End-to-end transparent
DS-Lite	Hybrid	Per-flow	Tunneling	IPv6 transport
MAP-E	Hybrid	Stateless	Tunneling	Encapsulation-based
MAP-T	Hybrid	Stateless	Translation	Translation-based
4+4	CPE	Per-allocation	Tunneling	Evolutionary transparency

Performance Analysis
Performance constitutes a critical evaluation criterion for addressing conservation mechanisms, affecting user experience, operational costs, and scalability.

State Table Scalability
The most significant performance challenge in stateful NAT deployments is the scalability of the translation state table. In CGNAT environments, the provider router must maintain state for potentially hundreds of thousands or millions of concurrent sessions across thousands of customers. This imposes substantial

Received*: 28-03-2026 -* ***Accepted****: 26-05-2026 -* ***Published****: 01-06-2026*

memory and CPU requirements. The scale of this challenge is substantial. Modern applications, even with multiplexed requests, can generate numerous concurrent sessions. Each session requires a state that must be created, maintained, aged, and eventually garbage-collected. In large-scale deployments, the state table can consume gigabytes of memory and require significant processing power for lookup and update operations. Stateless mechanisms offer a compelling alternative. A+P, by eliminating per-flow state, removes this scalability bottleneck entirely. The gateway simply needs to maintain static or quasi-static mappings between port ranges and customer identifiers, dramatically reducing state requirements.

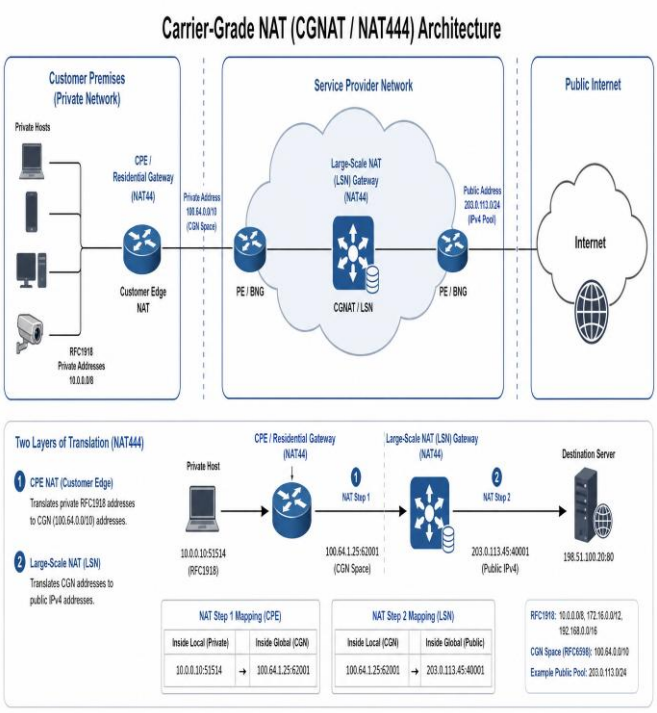

Figure 2. CGNAT Architecture

Latency Considerations

Each layer of address translation introduces processing latency. NAT translation overhead manifests as increased packet processing time, which translates to higher latency and potentially reduced throughput. Empirical studies have quantified these effects. Testing of NAT444 in large-scale operator environments has shown that the additional translation layer can increase latency by 20-30 milliseconds in 100,000-user scenarios. This latency is particularly impactful for real-time applications, including gaming, video conferencing, and voice over IP. The performance impact varies significantly across implementations. Hardware-accelerated and DPDK-based CGNAT implementations can achieve throughput of 10-100 Gbps with sub-10 microsecond latency. A DPDK-based CGNAT framework has demonstrated 15.5 times higher throughput than Linux kernel-based address translation services. These results indicate that while NAT imposes inherent overhead, optimized implementations can substantially mitigate performance degradation [11].

Throughput Considerations

Throughput in NAT systems is constrained by the packet processing pipeline. Each packet must undergo: 1. Header parsing and validation 2. State table lookup (for stateful NAT) 3. Address and port translation 4. Checksum recalculation 5. Packet forwarding.

The sequential nature of these operations, combined with the need for state synchronization in distributed deployments, creates throughput limitations. CGNAT deployments must carefully engineer their processing pipelines to avoid bottlenecks during peak traffic periods. Stateless mechanisms offer throughput advantages by eliminating the state lookup step. MAP-E and MAP-T, by employing deterministic mapping algorithms, can process packets with lower per-packet overhead. The A+P approach, similarly, avoids the per-flow state management that constrains stateful NAT throughput [12].

Impact of Double NAT (NAT444)

The NAT444 architecture introduces two layers of address translation: one at the customer premises (CPE NAT) and another at the carrier-grade NAT in the provider network. This double translation compounds performance challenges in several ways:

- **Increased latency**: Each translation layer adds processing delay
- **State multiplication**: State must be maintained at both layers

Received*: 28-03-2026 -* ***Accepted****: 26-05-2026 -* ***Published****: 01-06-2026*

- **Resource contention**: The CGNAT must serve all customers behind it
- **Failure amplification**: CGNAT failure affects all customers sharing that device

Testing by multiple service providers, including CableLabs, Time Warner Cable, and Rogers Communications, has documented the disruptive impact of this second NAT layer on common Internet applications.

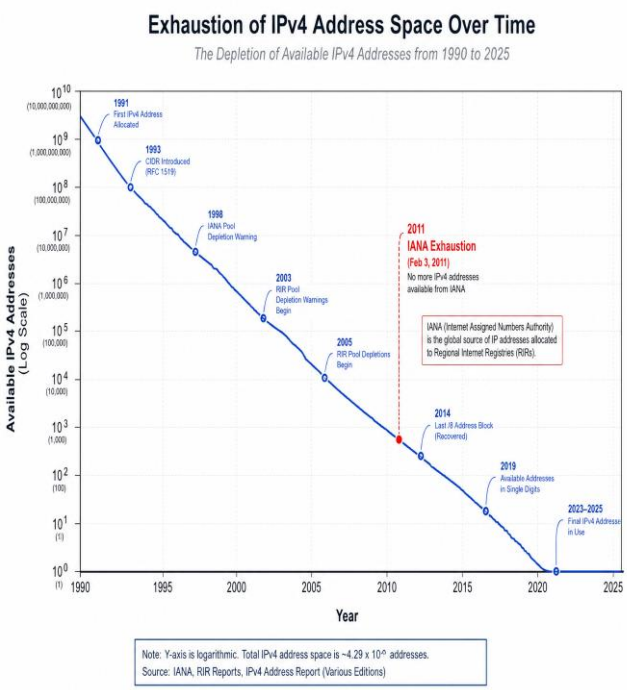

Figure 3. IPv4 Address Exhaustion Over Time

Transparency Implications

Transparency—the degree to which end hosts and applications can operate without awareness of or modification for network address translation—represents a critical evaluation dimension.

End-to-End Address Transparency

The most fundamental transparency issue concerns the visibility of IP addresses. In a transparent network, each host possesses a globally unique, routable IP address that is consistently visible to all communicating parties. NAT breaks this transparency by modifying source addresses as packets traverse network boundaries. As noted in RFC 2993, "NAT affects the transparency of end-to-end connectivity for transports relying on consistency of the IP header, and for protocols which carry that address information in places other than the IP header". This impacts:

- **IP-layer identification**: Hosts behind NAT cannot be uniquely identified by their IP address
- **Protocols with embedded addresses**: Protocols including FTP, SIP, and H.323 that carry IP addresses in application-layer payloads fail without Application-Level Gateways (ALGs)
- **Address-based authentication**: Security mechanisms relying on IP address verification are compromised

Application-Level Transparency

NAT deployment imposes significant challenges for application-layer protocols:

Peer-to-Peer Applications

P2P protocols rely on direct connections between peers. With CGNAT, such direct connections are often impossible, as all traffic must be routed through the carrier's NAT device. This can lead to increased latency, slower download speeds, and reduced overall performance. Studies of BitTorrent have shown that peers behind NAT tend not to receive favorable treatment, significantly decreasing download speeds and potentially degenerating the download into client-server interaction.

Real-Time Communications

Voice over IP (VoIP), video conferencing, and online gaming are particularly sensitive to NAT-induced issues. Session establishment, media traversal, and NAT keep-alive mechanisms introduce complexity and potential failure modes.

Received*: 28-03-2026 -* ***Accepted****: 26-05-2026 -* ***Published****: 01-06-2026*

Protocol Evolution
NAT has inhibited the development and deployment of new transport protocols. Protocols including DCCP, UDP-Lite, and SCTP face challenges in NAT environments, as they lack the widespread ALG support available for TCP and UDP.

Transparency Restoration Approaches
Several mechanisms aim to restore or preserve transparency:

4+4 Architecture
The 4+4 architecture leverages existing NATs and private address realms while enabling a return to end-to-end address transparency through incremental deployment. During transition, only NATs and end-hosts require updates, not network routers. Encapsulation serves as the primary tool for maintaining backward compatibility.

A+P Approach
By using port number bits as additional endpoint identifiers, A+P achieves equivalence to NAT with end-to-end transparency. The approach preserves the Internet's end-to-end principle to customer premises.

NAT Puncturing Techniques
For P2P communication, NAT punching enables the establishment of direct channels between peers behind NATs, given knowledge of the receiver's IP address. Testing across multiple carriers has demonstrated successful cross-connectivity in the majority of cases.

Architectural Implications for the End-to-End Principle
The deployment of NAT-based address conservation mechanisms carries profound architectural implications that extend beyond immediate operational concerns.

Fundamental Architectural Tension
The end-to-end principle, as articulated in the Internet's original design, places intelligence and state at the network edges. NAT inverts this model by introducing stateful processing in the network core. This represents not merely a practical compromise but a fundamental architectural shift. RFC 2993 documents the architectural consequences:

Single point of failure
NAT devices become critical network elements whose failure affects large populations.

ALG complexity
Application Layer Gateways introduce complexity and maintenance burden.

TCP state violations
NATs can violate TCP semantics.

Symmetric state management
Inconsistent state management across vendors creates interoperability issues.

Address collision risks
Private address reuse increases collision probability

Security Implications
The security implications of NAT are nuanced and often misunderstood:

Perceived Security Benefits:
NAT provides a form of network obscurity by hiding internal addressing structures. However, this is security through obscurity rather than true security, and should not be conflated with firewall functionality.

End-to-End Security
NAT breaks end-to-end security mechanisms. IPsec, in particular, faces challenges as NAT modifies IP headers and transport-layer checksums. While solutions including NAT-Traversal (NAT-T) exist, they add complexity and may compromise security properties.

Received*: 28-03-2026 -* ***Accepted****: 26-05-2026 -* ***Published****: 01-06-2026*

Lawful Intercept
NAT complicates lawful intercept and attribution, as multiple customers may share a single public IP address. This creates challenges for abuse attribution and forensic investigation.

Impact on Internet Evolution
The widespread deployment of NAT has altered the trajectory of Internet evolution:

Protocol Innovation
The need for NAT traversal has constrained protocol design. New protocols must either be NAT-friendly or incorporate traversal mechanisms, limiting architectural freedom.

IPv6 Transition
NAT has paradoxically slowed IPv6 adoption by providing a "good enough" solution to address exhaustion. While NAT extends IPv4 lifetime, it also reduces the urgency for IPv6 deployment, potentially prolonging the transition period.

Architectural Erosion
The End-Middle-End Research Group (EME) of the IRTF has specifically identified the progressive erosion of the end-to-end architecture, most notably by NATs and other middleboxes.

Stateless and Transparency-Preserving Alternatives
Stateless mechanisms offer a path toward reconciling address conservation with architectural principles:

A+P
By moving the address sharing function to the CPE and eliminating per-flow state in the network, A+P preserves the end-to-end model while achieving address conservation.

MAP-E and MAP-T
These stateless mechanisms enable address and port reuse without maintaining per-flow state in the network core. The deterministic mapping between customers and port ranges eliminates the need for stateful translation.

4+4 Architecture
This approach explicitly targets the restoration of end-to-end address transparency as deployment progresses. By requiring updates only to NATs and end-hosts, it provides an evolutionary path rather than a disruptive replacement.

Comparative Analysis and Trade-Off Discussion
The evaluation of NAT-based address conservation mechanisms reveals inherent trade-offs across performance, transparency, and architectural integrity dimensions.

Trade-Off Matrix
(Table 2) presents a comparative assessment of major mechanisms across key evaluation criteria.

Table 2. Comparative Analysis of NAT-Based Address Conservation Mechanisms

Mechanism	Address Conservation	Performance	Transparency	Architectural Compliance	Deployment Complexity
Traditional NAT	Moderate	Good	Low	Low	Low
CGNAT/NAT444	High	Moderate	Very Low	Very Low	Moderate
A+P	High	High	High	High	High
DS-Lite	High	Moderate	Moderate	Moderate	High
MAP-E	High	High	Moderate	Moderate	High
MAP-T	High	High	Moderate	Moderate	High
4+4	High	Moderate	High	High	Very High

Mechanism Selection Guidelines
The appropriate mechanism selection depends on deployment context and priorities:

For Immediate Address Conservation with Minimal Change
CGNAT/NAT444 offers the most direct path to address conservation with existing infrastructure. However, operators must accept performance degradation, transparency loss, and architectural compromise.

Received*: 28-03-2026 -* ***Accepted****: 26-05-2026 -* ***Published****: 01-06-2026*

For Preservation of End-to-End Transparency
A+P and 4+4 provide superior transparency at the cost of increased complexity and required host modifications. These approaches are most appropriate where end-to-end semantics are critical.

For IPv6 Transition Contexts
DS-Lite, MAP-E, and MAP-T are specifically designed for IPv6 transition scenarios. They enable IPv4-as-a-Service (IPv4aaS) while facilitating eventual IPv6-only operation.

The Path Forward
The continued reliance on NAT-based conservation mechanisms represents a transitional strategy rather than a permanent solution. The long-term resolution to address exhaustion remains IPv6 deployment, which provides sufficient address space to eliminate the need for address sharing.
However, the transition will be protracted. IPv4 reachability will remain necessary for the foreseeable future, and address sharing mechanisms will continue to play a role. The design of future mechanisms should prioritize:
1. Stateless operation to eliminate scalability bottlenecks
2. Transparency preservation to support application innovation
3. Architectural alignment with the end-to-end principle
4. Incremental deployability to facilitate gradual adoption

Conclusion
NAT-based address conservation mechanisms have served as a pragmatic response to IPv4 address exhaustion, enabling continued Internet growth in the face of fundamental addressing constraints. This paper has provided a comprehensive analysis of these mechanisms across the critical dimensions of performance, transparency, and architectural implications.
Our analysis reveals several key findings:
1. **Performance trade-offs are significant but manageable**: While stateful NAT introduces latency and scalability challenges, optimized implementations and stateless alternatives can substantially mitigate these issues.
2. **Transparency loss is inherent but not unavoidable**: Traditional NAT fundamentally breaks end-to-end address transparency, but mechanisms including A+P and 4+4 offer paths to transparency restoration.
3. **Architectural implications are profound**: NAT represents a fundamental departure from the Internet's end-to-end design philosophy, with consequences for security, protocol evolution, and network architecture.
4. **Stateless approaches offer a promising middle ground**: Mechanisms that eliminate per-flow state in the network preserve architectural principles while achieving address conservation.
The path forward requires balancing immediate operational needs against long-term architectural goals. While IPv6 remains the ultimate solution, the continued deployment of NAT-based mechanisms is inevitable during the transition period. Future research should focus on developing mechanisms that minimize performance impact, maximize transparency, and align with the Internet's fundamental architectural principles.

References
1. Egevang K, Francis P. The IP Network Address Translator (NAT). RFC 1631. Internet Engineering Task Force; 1994 May.
2. Hain T. Architectural Implications of NAT. RFC 2993. Internet Engineering Task Force; 2000 Nov.
3. Srisuresh P, Holdrege M. IP Network Address Translator (NAT) Terminology and Considerations. RFC 2663. Internet Engineering Task Force; 1999 Aug.
4. Rekhter Y, Moskowitz B, Karrenberg D, de Groot GJ, Lear E. Address Allocation for Private Internets. RFC 1918. Internet Engineering Task Force; 1996 Feb.
5. Jiang S, Casey T, Huang J, Zhang Y, Li Z. IPv4 Address Sharing Mechanism Classification and Tradeoff Analysis. IEEE/ACM Trans Netw. 2014 Apr;22(2):391-404.
6. Turányi Z, Valkó A, Campbell AT. 4+4: An Architecture for Evolving the Internet Address Space Back Toward Transparency. ACM SIGCOMM Comput Commun Rev. 2003 Jul;33(3):65-72.
7. Ohta M. Address plus Port (A+P) Approach to the IPv4 Address Shortage. IETF Internet-Draft. 2008.
8. Kanaris O, Pouwelse J. Mass Adoption of NATs: Survey and Experiments on Carrier-Grade NATs. arXiv:2311.04658. 2023.
9. Network Address Translation: Extending the Internet Address Space. IEEE Internet Comput. 2010 Jul-Aug;14(4):66-70.
10. Thomson M, Hu Z, Hain T. Assessing the Impact of Carrier-Grade NAT on Network Applications. RFC 7021. Internet Engineering Task Force; 2013 Sept.
11. Donley C, Grundemann C, Sarawat V, Sundaresan K, Gont F, Lear E, et al. NAT444 Impacts. IETF Internet-Draft. 2011.

Received*: 28-03-2026 -* ***Accepted****: 26-05-2026 -* ***Published****: 01-06-2026*