Internet Area Working Group
Internet-Draft
Intended status: Standards Track
Expires: 18 October 2026

M. Matolin
Valtrix Labs

16 April 2026

Global Anycast NAT64 Well-Known Prefix
draft-matolin-global-nat64-anycast-00

Abstract

This document defines a globally routable, anycast NAT64 service using the IPv6 prefix 2600:6464::/96 as a standardized translation substrate for IPv6-to-IPv4 connectivity.

The goal of this specification is to eliminate per-network NAT64 configuration complexity by introducing a single globally consistent NAT64 translation prefix operated as a distributed anycast service by participating Internet Service Providers, cloud providers, and content delivery networks.

The model assumes an IPv6-only client environment with mandatory IPv4 reachability via NAT64 translation. IPv4-only services remain reachable without modification.

IPv4 is not modified. IPv6 is not modified. Only translation placement and routing semantics are standardized.

This document defines:

* A globally shared NAT64 prefix (2600:6464::/96)
* Anycast-based NAT64 edge behavior
* Stateless IPv6-to-IPv4 synthesis rules
* Optional reverse mapping constraints (IPv4->IPv6 blocked)
* Operational requirements for participating networks

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Copyright Notice

Table of Contents

1. Introduction

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

NAT64 [RFC6146] enables IPv6-only clients to communicate with IPv4 servers via protocol translation. Current deployments require per-operator configuration of NAT64 prefixes, DNS64 behavior, and stateful translation pools.

This document proposes a globally standardized NAT64 prefix:

2600:6464::/96

and defines its use as a globally anycasted translation endpoint.

Instead of each operator deploying isolated NAT64 infrastructure, participating networks announce the prefix via BGP anycast, allowing the nearest translation edge to handle synthesis.

2. Motivation

Current NAT64 deployments suffer from:

* inconsistent prefix selection (multiple well-known prefixes)
* fragmented operational models
* lack of global routing consistency
* duplicated stateful NAT infrastructure per operator

This leads to:

* operational overhead
* inconsistent debugging models
* lack of deterministic routing behavior

A single global NAT64 anycast prefix provides:

* deterministic synthesis behavior
* reduced operational complexity
* unified debugging and telemetry model
* elimination of per-network NAT64 planning

3. Prefix Allocation

The IPv6 prefix 2600:6464::/96 is reserved as:

Global NAT64 Anycast Translation Prefix

Characteristics:

* MUST be routed via BGP anycast
* MUST NOT be subnetted beyond /96
* MUST NOT be used for non-translation purposes

* MUST be implemented by participating NAT64 edges

The last 32 bits represent the IPv4 address being synthesized.

Example:

2600:6464::0808:0808 -> 8.8.8.8

4. Anycast NAT64 Architecture

Participating operators advertise 2600:6464::/96 globally.

Routing behavior:

Client -> nearest NAT64 edge (anycast)
-> stateful or stateless translation
-> IPv4 Internet

The architecture is intentionally stateless at routing level and stateful at translation edge only.

All NAT state is local to the terminating edge.

5. IPv6-to-IPv4 Mapping

Mapping rule:

IPv6 address = 2600:6464:0:0:W.X.Y.Z

Where:

* W.X.Y.Z is IPv4 destination
* NAT64 edge extracts IPv4 literal
* packet is translated and forwarded

No DNS dependency is required if literal IPv4 embedding is used.

6. DNS64 Interaction

DNS64 MAY synthesize AAAA records using prefix 2600:6464::/96.

Example:

A record: example.com -> 93.184.216.34

Synthesized AAAA:
2600:6464::5db8:d822

DNSSEC considerations:

* synthesis occurs after validation
* no modification to authoritative DNS required

7. Reverse Traffic Policy

IPv4 -> IPv6 direct access via NAT64 prefix is explicitly disallowed.

Reason:

* prevents external injection into IPv6 space
* preserves IPv6-only security boundary model
* avoids bidirectional NAT ambiguity

Only IPv6-originated sessions MAY traverse NAT64 edges.

8. Operational Requirements

Any network participating in global NAT64 anycast MUST:

* announce 2600:6464::/96 via BGP
* implement IPv4 translation backend
* support at least 10M concurrent NAT sessions
* provide per-edge state isolation
* implement deterministic failover within anycast cluster

SHOULD:

* log translation mappings (privacy-aware)
* support stateless mapping mode for CDN workloads
* expose telemetry via standard API

9. Security Considerations

Threat model includes:

* NAT64 edge spoofing
* prefix hijacking via BGP
* session state exhaustion attacks

Mitigations:

* RPKI validation for prefix 2600:6464::/96
* strict ingress filtering at NAT edges
* per-source rate limiting
* mandatory DNS64 validation where applicable

Anycast NAT64 edges MUST NOT forward packets without valid IPv4 extraction context.

10. IANA Considerations

IANA is requested to:

* reserve 2600:6464::/96 as "Well-Known Global NAT64 Anycast Prefix"
* register NAT64 anycast service type identifier
* document routing constraints for global translation services

11. Acknowledgements

Inspired by decades of NAT64 deployments, CGNAT scaling pain, and the universal human desire to stop dealing with IPv4.

12. References

12.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, April 2011, <https://www.rfc-editor.org/info/rfc6146>.

12.2. Informative References

Author’s Address

M. Matolin
Valtrix Labs
Email: meizfl@valtrix.org