IPv6 Operations
Internet-Draft
Intended status: Standards Track
Expires: 17 November 2026

W. Kumari
J. Linkova
Google, LLC

16 May 2026

NAT64 WKP
draft-ietf-v6ops-nat64-wkp-1918-02

Abstract

This document removes the requirement introduced in Section 3.1 of RFC6052 that the NAT64 Well-Known Prefix 64:FF9B::/96 MUST NOT be used to represent non-global IPv4 addresses, such as those defined in [RFC1918] or listed in Section 3 of [RFC5735]. The proposed change enables IPv6-only nodes to reach IPv4-only services with non-global addresses by leveraging the Well-Known Prefix.

About This Document

This note is to be removed before publishing as an RFC.

Status information for this document may be found at https://datatracker.ietf.org/doc/draft-ietf-v6ops-nat64-wkp-1918/.

Discussion of this document takes place on the IPv6 Operations Working Group mailing list (mailto:v6ops@ietf.org), which is archived at https://mailarchive.ietf.org/arch/browse/v6ops/. Subscribe at https://www.ietf.org/mailman/listinfo/v6ops/.

Source for this draft and an issue tracker can be found at https://github.com/furry13/6052-update-wkp1918.

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/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 17 November 2026.

Table of Contents

1. Introduction

Section 3.1 of [RFC6052] prohibits IPv4/IPv6 translators from using the Well-Known Prefix (WKP, 64:FF9B::/96) to represent non-global IPv4 addresses, such as those defined in [RFC1918] or listed in Section 3 of [RFC5735].

This restriction is relatively straightforward to implement in DNS64 [RFC6147]: a DNS64 server simply avoids synthesizing an AAAA record using the WKP if the original A record contains a non-global IPv4 address. However, this requirement introduces significant operational challenges for systems that do not rely on DNS64 and instead use local synthesis such as CLAT (Customer-side Translator, [RFC6877]), or similar approaches.

Enterprise and other closed networks often require IPv6-only nodes to communicate with both internal (e.g., using RFC1918 addresses) and external (Internet) IPv4-only destinations. The restriction in Section 3.1 of RFC6052 prevents such networks from utilizing the WKP and, consequently, from relying on public DNS64 servers (e.g. forwarding requests for external zones to public DNS64) which utilize the WKP in order to maximize compatibility.

Using two NAT64 prefixes — the WKP for Internet destinations and a Network-Specific Prefix (NSP) for non-global IPv4 addresses — is not a feasible solution for nodes performing local synthesis or running CLAT. None of the widely deployed NAT64 Prefix Discovery mechanisms ([RFC7050], [RFC8781]) provide a method to map a specific NAT64 prefix to a subset of IPv4 addresses for which it should be used.

According to Section 3 of [RFC7050], a node must use all learned prefixes when performing local IPv6 address synthesis. Consequently, if a node discovers both the WKP and the NSP, it will use both prefixes to represent global IPv4 addresses. This duplication significantly complicates security policies, troubleshooting, and other operational aspects of the network.

Prohibiting the WKP from representing non-global IPv4 addresses offers no substantial benefit to IPv6-only or IPv6-mostly deployments. Simultaneously, it substantially complicates network design and the behavior of nodes.

Given the recent operational experience in deploying IPv6-only and IPv6-mostly networks, it is desirable to allow translators to use a single prefix (including the WKP) to represent all IPv4 addresses, regardless of their global or non-global status. This simplification would greatly improve the utility of the WKP in enterprise networks.

2. Conventions and Definitions

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.

2.1. Terminology

This document reuses the Terminology section of [RFC6052].

3. RFC6052 Update

This document updates Section 3.1 of [RFC6052] ("Restrictions on the Use of the Well-Known Prefix") as follows:

OLD TEXT:

===

The Well-Known Prefix MUST NOT be used to represent non-global IPv4 addresses, such as those defined in [RFC1918] or listed in Section 3 of [RFC5735]. Address translators MUST NOT translate packets in which an address is composed of the Well-Known Prefix and a non-global IPv4 address; they MUST drop these packets.

===

NEW TEXT:

===

The Well-Known Prefix MAY be used to represent non-global IPv4 addresses, such as those defined in [RFC1918] or listed in Section 3 of [RFC5735].

Unmanaged client-side translators (CLATs) MUST translate packets in which an address is composed of the Well-Known Prefix and a non-global IPv4 address by default.

Provider-side translators (PLATs) MUST translate such packets unless configured otherwise. Because administrators may rely on dropping these packets as an implicit security policy, PLAT implementations MAY choose not to translate such packets by default. However, such PLAT implementations SHOULD provide a configuration knob to enable translation for these packets.

===

As noted in Errata 5547 ([EID5547]):

IPv4 packets with private destination addresses are routinely translated to IPv4 packets with global destination addresses in NAT44. Similarly, an IPv6 packet with a destination address representing a private IPv4 address [RFC6052] can be translated to an

IPv4 packet with a global destination address by NAT64 [RFC6146]. If a 464XLAT CLAT cannot translate a private IPv4 address to an IPv6 address using the NAT64 /96 prefix and that IPv4 address [RFC6052], then the packet may not be translated to an IPv4 packet with a global address by the 464XLAT PLAT (stateful NAT64). This changes the intent of the sender, and in so doing violates the end to end principle.

Removing the requirement introduced in RFC 6052 Section 3.1 addresses this errata.

4. Operational Considerations

There may be cases when it is desirable to ignore translation of private use IPv4 addressing due to internal policy or overlapping internal networks. It is important to note, however, that overlapping networks in IPv6 translated addresses are also overlapping in IPv4, and so behavior will be similar across protocols in the vast majority of use cases. Environments reliant on [RFC7050] may be required to create configurations which address the filtering of private use IPv4 addressing if there is an expectation of compliance with the original section 3.1.

4.1. Existing Behavior

Testing and operational experience with existing CLAT implementations (both mobile and non-mobile) have revealed highly inconsistent behavior regarding the original restriction in Section 3.1 of [RFC6052]. While some implementations strictly comply with the original requirement and drop packets destined for non-global IPv4 addresses, many other widely deployed CLATs completely ignore this restriction and translate the packets.

This inconsistency creates significant operational challenges. Network operators are unable to predictably determine how unmanaged, client-side devices will handle traffic directed to internal IPv4 services. This unpredictable dropping or translating of packets on the client side severely complicates network design, security policies, and troubleshooting.

By formalizing the requirement that unmanaged CLAT implementations MUST translate these packets by default (as updated in Section 3), and allowing PLAT devices to translate these packets, this document provides clear, standardized instructions to implementers. This resolves the current operational ambiguity, ensuring predictable behavior across all client ecosystems and aligning the standard with the practical realities of modern IPv6-mostly and IPv6-only deployments.

Furthermore, where client-side translation and local synthesis are used, it is currently not feasible to employ more than one translation prefix, especially if different prefixes must be used for different IPv4 destinations. None of the widely deployed NAT64 Prefix Discovery mechanisms ([RFC7050], [RFC8781]) provide a method to map a specific NAT64 prefix to a subset of IPv4 addresses for which it should be used.

4.2. Use of Network Specific Prefix

Use of a network specific prefix such as provided by [RFC8215] does not preclude the removal of section 3.1 as a MUST requirement. If a network employs a network specific prefix the behavior of synthesizing a private use IPv4 address is not prevented by standard. The use of a network specific prefix implies the existence of a local mechanism for synthesizing IPv6 addresses based on that specific prefix, and thereby rules out use of a public DNS64 resolver in the vast majority of cases, as large scale public DNS64 resolvers use the WKP to maximize compatibility.

5. Security Considerations

Legitimizing packets where the IPv6 destination address is composed of the WKP and a non-global IPv4 address does not, inherently, introduce new security considerations. Whether a specific traffic flow between an IPv6-only source and a non-global IPv4 destination (or any flow to a non-global IPv4 destination) is legitimate is a matter of local network topology and administrative policy. However, existing NAT64 implementations compliant with RFC 6052 are expected to drop such packets. Administrators may be relying on this implicit filtering as a built-in security mechanism to prevent unauthorized access to private IPv4 infrastructure, rather than implementing explicit security policies. This reliance is particularly prevalent in managed NAT64 (PLAT) environments.

Modifying the recommended behavior to allow such address compositions may, in the absence of explicit filtering, enable traffic flows that were previously prohibited by the translator’s default logic. To mitigate this risk, existing managed NAT64 implementations compliant with RFC 6052 SHOULD NOT alter their default dropping behavior. Instead, they SHOULD provide a configuration knob to enable this functionality, ensuring that the transition to supporting non-global addresses is an intentional administrative action accompanied by a review of local security policies.

Furthermore, administrators should not rely on the internal verification logic of the translator to enforce security boundaries. Instead, explicit policies such as access control lists (ACLs), firewall policies or NAT rules must be used to define authorized traffic patterns through the translator.

6. IANA Considerations

This document has no IANA actions.

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/rfc/rfc2119>.

[RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052, DOI 10.17487/RFC6052, October 2010, <https://www.rfc-editor.org/rfc/rfc6052>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/rfc/rfc8174>.

7.2. Informative References

[EID5547] "Errata ID 5547: NAT64 Well-Known Prefix SHOULD NOT be used for Private Use IPv4 Addresses", n.d., <https://www.rfc-editor.org/errata/eid5547>.

[RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. J., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <https://www.rfc-editor.org/rfc/rfc1918>.

[RFC5735] Cotton, M. and L. Vegoda, "Special Use IPv4 Addresses", RFC 5735, DOI 10.17487/RFC5735, January 2010, <https://www.rfc-editor.org/rfc/rfc5735>.

[RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, DOI 10.17487/RFC6146, April 2011, <https://www.rfc-editor.org/rfc/rfc6146>.

[RFC6147] Bagnulo, M., Sullivan, A., Matthews, P., and I. van Beijnum, "DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", RFC 6147, DOI 10.17487/RFC6147, April 2011, <https://www.rfc-editor.org/rfc/rfc6147>.

[RFC6877] Mawatari, M., Kawashima, M., and C. Byrne, "464XLAT: Combination of Stateful and Stateless Translation", RFC 6877, DOI 10.17487/RFC6877, April 2013, <https://www.rfc-editor.org/rfc/rfc6877>.

[RFC7050] Savolainen, T., Korhonen, J., and D. Wing, "Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis", RFC 7050, DOI 10.17487/RFC7050, November 2013, <https://www.rfc-editor.org/rfc/rfc7050>.

[RFC8215] Anderson, T., "Local-Use IPv4/IPv6 Translation Prefix", RFC 8215, DOI 10.17487/RFC8215, August 2017, <https://www.rfc-editor.org/rfc/rfc8215>.

[RFC8781] Colitti, L. and J. Linkova, "Discovering PREF64 in Router Advertisements", RFC 8781, DOI 10.17487/RFC8781, April 2020, <https://www.rfc-editor.org/rfc/rfc8781>.

Acknowledgments

The authors would like to thank Mohamed Boucadair, Nick Buraglio, Lorenzo Colitti, Suresh Krishnan, Ted Lemon, Jordi Palet for their helpful comments and suggestions on this document.

Appendix: Example flow

{ *Ed note*: Nick Buraglio has suggested that we include an example flow here. I think that this can be removed before publication, but it might be helpful to include fur discussion / during LC, etc}

To illustrate the updated normative behavior, consider an IPv6-only network utilizing 464XLAT [RFC6877] where an administrator wishes to provide access to an internal, IPv4-only corporate service hosted at 10.1.2.3.

Scenario A: Unmanaged CLAT to Managed PLAT Flow

An IPv4-only application on an unmanaged client device generates an IPv4 packet destined for 10.1.2.3.

The local CLAT intercepts the IPv4 packet and synthesizes an IPv6 destination address by prepending the Well-Known Prefix: 64:ff9b::10.1.2.3.

CLAT Behavior: Under the updated guidance in Section 3, the CLAT MUST translate this packet by default, ignoring the non-global nature of the embedded IPv4 address, and forward the resulting IPv6 packet to the network.

The IPv6 network routes the packet to the managed PLAT (NAT64 gateway).

PLAT Behavior: Upon receiving the packet destined for 64:ff9b::10.1.2.3, the PLAT evaluates its local configuration:

Permit: If the administrator has explicitly enabled translation for non-global addresses (or left the default translation behavior enabled), the PLAT translates the packet back to IPv4 and forwards it to 10.1.2.3.

Drop: If the administrator relies on a default-drop posture for non-global addresses or has explicitly configured an access control list (ACL) blocking this range, the PLAT drops the packet.

Scenario B: Native IPv6 Host to Managed PLAT

An IPv6-capable host (without a local CLAT) needs to communicate with the same internal service. It acquires the destination address 64:ff9b::10.1.2.3 (e.g., via DNS64, local synthesis, or explicit application configuration).

The host transmits the IPv6 packet, which is routed to the PLAT.

PLAT Behavior: The PLAT applies the same configuration logic as in Scenario A. It MUST translate the packet to IPv4 and forward it to 10.1.2.3 unless local administrative policy configures it to drop the packet.

Authors' Addresses

Warren Kumari
Google, LLC
Email: warren@kumari.net

Jen Linkova
Google, LLC
Email: furry13@gmail.com