Workgroup: IPv6 operations
Internet-Draft: draft-link-v6ops-6mops
Published: 4 March 2024
Intended Status: Informational
Expires: 5 September 2024
Authors: J. Linkova
Google

IPv6-Mostly Networks: Deployment and Operations Considerations

Abstract

This document discusses an deployment scenario called "an IPv6-Mostly network", when IPv6-only and IPv4-enabled endpoints coexist on the same network (network segment, VLAN, SSID etc).

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 5 September 2024.

Table of Contents

1. Introduction

While most network operators initially deploy IPv6 alongside their existing IPv4 infrastructure, pure IPv6-only networks remain uncommon outside of the mobile carrier space. This dual-stack approach is seen as a necessary transition phase, allowing operators to gain experience with IPv6 while minimizing disruption.

However, dual-stack networks don't address the core problem driving IPv6 adoption: IPv4 address exhaustion. They still require the same amount of IPv4 resources as IPv4-only networks. Even worse, this dual-stack approach often becomes a long-term crutch. Many applications still rely on IPv4, creating a chicken-and-egg problem: IPv6-only networks seem impractical with so many incompatible applications, yet applications continue to rely on IPv4 because IPv6-only networks are rare.

The less control a network operator has over devices and applications, the more difficult it is to break IPv4 dependencies and move to IPv6-only. This is particularly challenging in enterprise networks with legacy IPv4-dependent applications and public WiFi networks where operators cannot guarantee device compatibility.

To enable a gradual migration, operators need to identify which devices can function in IPv6-only mode and which cannot. Creating separate network segments for each type introduces complexity and scalability issues – a major hurdle to IPv6-only adoption.

A more desirable approach is to deploy so-called "IPv6-mostly" network that provides IPv4 on demand. This allows IPv6-capable devices to remain IPv6-only while seamlessly supplying IPv4 to those that require it.

This document explores the requirements, recommendations, and challenges associated with deploying IPv6-mostly networks in enterprise and public WiFi environments. While the principles discussed may be applicable to other network types, this document's focus remains on these specific use cases.

2. Requirements Language

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.

3. Terminology

This document reuses most of Terminology section from [RFC8925].

Endpoint: A device connected to a network and considered a host From the operator's perspective. However, some endpoint can also extend

the network to other physical or logical systems, thereby assuming routing functions. Examples include:

* Corporate laptop: While primarily a host, it might run virtual systems and route traffic to them, extending the network and acting as a router.

* Mobile phone with tethering enabled: Acts as a host on the WiFi network, but also as a router for tethered devices, potentially without the operator's knowledge or consent.

Network segment: a link (VLAN, a broadcast domain etc) where hosts share the same IP subnet.

4. Solution Overview

In a nutshell, an IPv6-mostly network is very much similar to a dual-stack one with two additional key elements:

* The network provides NAT64 ([RFC6146]) functionality ([RFC6146]), enabling IPv6-only clients to communicate with IPv4-only destinations.

* The DHCPv4 infrastructure processing DHCPv4 Option 108 as per [RFC8925].

Upon connecting to an IPv6-mostly network segment, an endpoint configures its IP stack based on its capabilities:

* IPv4-Only Endpoint: Acquires an IPv4 address through DHCPv4.

* Dual-Stack Endpoint (Not IPv6-Only Capable): Configures IPv6 addresses via Stateless Address Autoconfiguration (SLAAC) and, optionally, DHCPv6. Additionally, it obtains an IPv4 address via DHCPv4.

* IPv6-only capable endpoint configures its IPv6 addresses and, while performing DHCPv4, includes option 108 ([RFC8925]) into the Parameter Request List. The DHCP server returns the option and, as per [RFC8925], the endpoint forgoes requesting an IPv4 address, remaining in IPv6-only mode.

An IPv6-mostly network segment can support a mix of IPv4-only, dual-stack, and IPv6-only devices. IPv6-only endpoints utilize the network-provided NAT64 to reach IPv4-only destinations.

The following sections discussed the various solution elements in more details.

4.1. IPv6-Only Capable Endpoints

The term "IPv6-only capable endpoint" lacks a strict technical definition. It broadly describes a device that can function without native IPv4 connectivity/IPv4 addresses, providing the same user experience. The most common way to achieve this is by implementing a customer-side translator (CLAT) as specified in 464XLAT ([RFC6877]). Devices which support CLAT (such as mobile phones) are known to operate without issues in IPv6-only mode. In some cases, however, a network administrator may consider a device IPv6-only capable even without CLAT implementation. For example, if all applications run on the device have been tested and confirmed to operate in NAT64 envinronment without any IPv4 dependencies.

4.2. IPv6-Only and IPv4-enabled Endpoints Coexistence

One effective way to restrict IPv4 addresses solely to devices that require them is to enable support for the IPv6-Only Preferred DHCPv4 option (Option 108, [RFC8925]) on the network's DHCP infrastructure. Most CLAT-enabled systems also support Option 108. By recognizing this option, the network can configure those devices as IPv6-only, allowing them to use CLAT fo providing IPv4 address to the local endpoint's network stack.

Certain devices, such as resource-constrained embedded systems, may operate in IPv6-only mode without CLAT if their communication is limited to IPv6-enabled destinations. Since these systems often lack Option 108 support, administrators may need alternative methods to prevent IPv4 address assignment. One approach is to block IPv4 traffic at the switchport level. This could involve:

*Static ACL: Applying a static filter with a "deny ip any any" rule.

*Dynamic ACL via RADIUS: If 802.1x authentication is in use, RADIUS can provide an ACL blocking all IPv4 traffic.

The ACL-based approach has some scalability implications and increases operational complexity, therefore it could only be recommended as a stopgap solution.

4.3. Access to IPv4-only Destinations

4.3.1. NAT64

IPv6-only endpoints require NAT64 to access IPv4-only destinations. Quite often operators choose to combine NAT44 and NAT64 functions. However, if not all internal services are IPv6-enabled, then NAT64 might need to be performed closer to the clients. If IPv4-only internal destinations are using [RFC1918] address space, then the

operator MUST NOT use the well-known prefix 64:ff9b::/96 for NAT64 (see section 3.1 of [RFC6052]).

4.3.2. 464XLAT

Enabling CLAT (Customer-side translator) on endpoints is essential for seamless operation of IPv4-only applications in IPv6-only environments. CLAT provides an [RFC1918] address and IPv4 default route, ensuring functionality even without a native IPv4 address from the network. Without CLAT, IPv4-only applications would fail, negatively impacting user experience and increasing support overhead.

Recommendations for Network Administrators controlling the endpoints:

- CLAT + DHCPv4 Option 108: If the network administrator can control endpoint configuration, CLAT SHOULD be enabled on endpoints sending DHCPv4 Option 108. This streamlines the transition.

- Option 108 Without CLAT MAY be enabled if the administrator is willing to identify and fixIPv4-only systems/applications, or if all applications are confirmed to work in IPv6-only mode.

4.3.3. Signalling NAT64 Prefix to Hosts

Hosts running 464XLAT need to discover the PREF64 (the IPv6 prefix used by NAT64). The network administrator SHOULD configure the first-hop routers to include PREF64 information in Router Advertisements as per ([RFC8781]) even if the network provides DNS64 (so hosts can use DNS64-based prefix discovery, [RFC7050]). This is required as hosts or individual applications might have custom DNS configuration (or even run a local DNS server) and ignore DNS64 information provided by the network, so they can not use [RFC7050] method to detect PREF64. In the absense of PREF64 information in Router Advertisements such systems would not be able to run clat, which would cause connectivity issues for all IPv4-only applications running on the affected device. As such device wouldn't be able to use the network-provided DNS64, access to IPv4-only destination would be impacted as well. At the time of writing all major OSes supporting DHCPv4 option 108 and enabling clat automatically also support [RFC8781]. Therefore providing PREF64 information in RAs can reliably mitigate the impact of custom DNS configuration on those systems.

Receiving PREF64 information in RAs also speeds up the clat start up time, so an IPv4 address and default route become available for applications much faster.

4.3.4. DNS vs DNS64

While DNS64 (with NAT64) enables IPv6-only endpoints to access IPv4-only destinations, it has several drawbacks:

- DNSSEC Incompatibility: DNS64 responses fail DNSSEC validation.
- Custom Resolvers: Endpoints or applications configured with custom resolvers can not benefit from the DNS64 provided by the network.
- Additional requirements for application: to benefit from DNS64, applications need to be IPv6-enabled, use DNS (do not use IPv4 literals). Many applications do not satisfy those requirements and therefore fail if the endpoint does not have an IPv4 address/native IPv4 connectivity.

If the network provides PREF64 in RAs (Section 4.3.3) and all endpoints are guaranteed to have CLAT enabled, DNS64 is unnecessary and SHOULD NOT be enabled. However, if some IPv6-only devices might lack CLAT support, the network MUST provide DNS64 unless those endpoints are guaranteed never to need IPv4-only destinations (for example, in case of a specialized network segment communicating solely with IPv6-capable destinations).

5. Solution Benefits

5.1. Benefits Compared to Dual-Stack

IPv6-Mostly networks offer significant advantages over traditional dual-stack models where endpoints have both IPv4 and IPv6 addresses:

- Drastically Reduced IPv4 Consumption: Dual-stack deployments don't solve the core problem of IPv4 address exhaustion. IPv6-Mostly allows to significantly reduce IPv4 consumption, as well as to reclaim IPv4 space. This reduction depends on endpoint capabilities (DHCPv4 Option 108 and CLAT support). In real-world scenarios, like conference Wi-Fi, 60-70% of endpoints may support IPv6-only operation, potentially allowing 2-4 times smaller IPv4 subnets.
- Simplified Operations: Managing dual-stack networks means running two network planes simultaneously, increasing complexity, costs, and the potential for errors. IPv6-Mostly allows to phase out IPv4 from many endpoints, streamlining operations and improving overall network reliability.
- Reduced Dependency on DHCPv4: As more devices operate seamlessly in IPv6-only mode, the criticality of DHCPv4 service diminishes significantly. This allows operators to scale down DHCPv4

infrastructure or operate it with less stringent SLOs, optimizing costs and resource allocation.

5.2. Benefits Compared to a Dedicated IPv6-Only Network

Traditional IPv6-only adoption involved separate networks alongside dual-stack ones. IPv6-Mostly offers significant improvements:

* Enhanced Scalability: Separate IPv6-only networks double the number of SSIDs in wireless environments, causing channel congestion and degrading performance. IPv6-Mostly doesn't require additional SSIDs. Similarly, it allows IPv4 and IPv6-only devices to coexist on the same wired VLANs, eliminating the need of additional VLANs.

* Operational Simplicity: Managing one network segment for all clients (regardless of IPv4 needs) simplifies operations, improves user experience (no more confusing SSID choices), and reduces support tickets related to mismatched connections. Dynamic VLAN assignment becomes easier without device-specific IPv6 capability tracking.

* Optimized IPv4 Consumption: User-selected dual-stack networks often lead to unnecessary IPv4 use, as users often connect IPv6-only capable devices to a dual-stack network. IPv6-Mostly network allocates IPv4 addresses only when devices don't advertise IPv6-only capability (DHCPv4 Option 108).

* Improved Problem Visibility: User-selected fallback to dual-stack networks can mask issues with IPv6-only operation, hindering problem reporting and resolution. IPv6-Mostly forces users to work through any issues, improving identification and enabling fixes for smoother long-term migration.

* Flexible, Incremental Migration: IPv6-Mostly allows for gradual device migration on a per-segment basis. Devices become IPv6-only only when deemed fully compatible with that mode.

6. Incremental Rollout Considerations

Migrating endpoints to IPv6-only fundamentally changes network dynamics by removing the IPv4 safety net. This includes the masking effect of Happy Eyeballs. IPv6 connectivity issues become far more prominent, including those previously hidden within dual-stack environments. Operators should be prepared to discover and troubleshoot issues in both endpoints and network infrastructure, even if the dual-stack network appeared problem-free.

Some rollout considerations are discussed in the following sections.

6.1. Per-Device and Per-Subnet Incremental Rollout

Limited control over endpoint configuration necessitates a per-subnet rollout, incrementally enabling Option 108 processing in DHCP. If endpoint control exists, per-device rollout is possible (at least for OSes with configurable Option 108). Note that some OSes unconditionally enable Option 108 support, becoming IPv6-only the moment it's activated on the server side. The following approach is RECOMMENDED:

- DHCP Server-Side Activation: Enable Option 108 processing. Some OSes automatically switch to IPv6-only. Rollback at this stage affects the entire subnet.

- Controlled Endpoint Activation: Enable Option 108 on managed endpoints with per-device rollback possible.

6.2. Rollback Approach

For quick rollback, the administrator SHOULD start with a minimal Option 108 value (300 seconds, Section 3.4 of [RFC8925]) and increase this value as the IPv6-mostly network proves reliable, reducing the likelihood of full-scale rollback.

6.3. Opt-In and Opt-Out Modes

Before user-facing deployment, the administrator SHOULD consider a dedicated IPv6-mostly proof-of-concept network for early adopters. While this temporarily sacrifices some IPv6-mostly benefits (Section 5.2), it provides valuable operational experience and early issue detection.

- Opt-In Phase: Invite tech-savvy early adopters to enable Option 108 and report issues. While response rates may be low, dedicated participants provide valuable troubleshooting data.

- Opt-Out Phase: Incrementally enable Option 108. Allow selective disabling for problematic endpoints. Disabling Option 108 SHOULD NOT be possible without a problem reports to ensure issue tracking and resolution.

In some scenarios (see Section 7.3.1) the administrator MAY keep a dual-stack network as a last resort fallback mechanism but SHOULD prevent usesrs from connecting to it accidentially (e.g. it should be a hidden protected SSID).

7. Operational Considerations

7.1. Address Assignment Policy

As outlined in Section 6.3 of [RFC6877], CLAT requires either a dedicated IPv6 prefix or, if unavailable, a dedicated IPv6 address. Currently (2024), all implementations use SLAAC for CLAT address acquisition. Therefore, to enable CLAT functionality within IPv6-mostly network segments, first-hop routers must advertise a Prefix Information Option (PIO) containing a globally routable SLAAC-suitable prefix with the 'Autonomous Address-Configuration' (A) flag set to zero.

7.2. Extension Headers

Being an IPv6-specific concept, IPv6 extension headers are often neglected or even explictly prohibited by security policies in dual-stack networks. The issues caused by blockig extension headers might be masked by the presense of Happy Eyeballs but become highly visible when there is no IPv4 to fallback to.

The network SHOULD permit at least the following extension headers:

* Fragment Header (Section 4.5 of [RFC8200]). Section 7.3.5 discusses the fragmentation in more details.

* ESP Header, which is used for IPSec traffic, such as VPN and WiFi Calling.

7.3. Typical Issues

IPv6-mostly networks expose hidden issues by removing the IPv4 safety net. While implementation bugs vary greatly and are beyond the scope of this document, this section focus on common problems caused by configuration, topology, or design choices. It's crucial to note that these issues likely pre-exist in dual-stack networks, but remain unnoticed due to IPv4 fallback.

7.3.1. Hosts with Disabled or Disfunctional IPv6

Historically, tech support often advised disabling IPv6 as a quick workaround, leading to devices with disabled IPv6. Similarly, corporate IT may have disabled or filtered IPv6 under the assumption that it's not widely used. Such endpoints requesting Option 108 will fail to connect in an IPv6-mostly network, as they won't receive IPv4 addresses and Ipv6 is disabled.

Administrators controlling endpoints SHOULD ensure those endpoints have IPv6 enabled and operational before migrating the network to IPv6-mostly mode.

7.3.2. Network Extension

IPv4's NAT44 allows endpoints to extend connectivity to downstream systems without upstream network awareness or permission. This creates challenges in IPv6-mostly deployments where endpoints lack IPv4 addresses:

Solutions and trade-offs:

- Using DHCPv6-PD to allocate prefixes to endpoints ([I-D.ietf-v6ops-dhcp-pd-per-device]). Provides downstream systems with IPv6 addresses and native connectivity.

- Enabling the CLAT function on the endpoint. This scenario is similar to the Wireline Network Architecture described in Section 4.1 of [RFC6877]. The downstream systems would receive IPv4 addresses and their IPv4 traffic would be translated to IPv6 by the endpoint. However thios approach leads to the downstream systems using IPv4 only and not benefiting from end2end IPv6 connectivity. To enable IPv6 benefits, combine this with IPv6 Prefix Delegation as above.

- Bridging and ND Proxy: The endpoint bridges IPv6 traffic and masks downstream devices behind its MAC address. This can lead to scalability issues (Section 7.3.3) due to the single MAC being mapped to many IPv6 addresses.

7.3.3. Multiple Addresses per Device

Unlike IPv4, where endpoints typically have a single IPv4 address per interface, IPv6 endpoints inherently use multiple addresses:

- Link-local address.

- Temporary address (common on mobile devices for privacy)

- Stable address (for long-term identification)

- CLAT address (in IPv6-mostly/IPv6-only networks)

Endpoints with containers, namespaces, or ND proxy functions can have even more addresses. This poses challenges for network infrastructure devices (SAVI switches, wireless access points etc) that map MAC addresses to IPv6 addresses, often with limits to prevent resource exhaustion or DoS attacks. When the number of IPs per MAC limit is exceeded, infastructure devices behavior varies across implementations, leading to inconsistent connectivity loss: while some systems drop new addresses, others delete older entries, causing previously functional addresses to lose connectivity. In all

those cases Endpoints and applications don't receive explicit signalling about the address becoming unusable.

While allocating prefixes to endpoints via DHCP-PD ([I-D.ietf-v6ops-dhcp-pd-per-device]) alows to eliminate the issue and corresponding scalability concerns, that solution migth not be supported by all endpoints. The network administrator SHOULD ensure that the deployed network infastructure devices allow sufficient number of IPv6 addresses to be mapped to a client's MAC and SHOULD monitor for events, indicating that the limit has been reached (such as syslog messages etc).

7.3.4. Host Mobility and Renumbering

Networks employing dynamic VLAN assignment (e.g., based on 802.1x or MAC-based authentication) can cause endpoints to move between VLANs and IPv6 subnets. As client operating systems do not always handle changes in link-layer state (e.g., VLAN changes) correctly, this mobility often leads to inconsistent IP stack behavior on operating systems, resulting in the persistence of old subnet addresses and potential connectivity issues due to incorrect source address selection. [I-D.link-v6ops-gulla] provides further analysis and potential solutions.

7.3.5. Fragmentation

As the basic IPv6 header is 20 bytes longer than the IPv4 header, transating from IPv4 to IPv6 can result in packets exceeding the path MTU on the IPv6 side. In that case NAT64 creates IPv6 packets with the Fragment Header (see Section 4 of [RFC6145] for more details. As per [RFC6145], by default the translator fragments IPv4 packets so that they fit in 1280-byte IPv6 packets. It means that all IPv4 packets larger than 1260 bytes are fragmented (or dropped if the DF bit is set).

Administrators SHOULD maximize the path MTU on the IPv6 side (from the translator to IPv6-only hosts) to minimize fragmentation. NAT64 devices SHOULD be configured to use the actual path MTU on the IPv6 side when fragmenting IPv4 packets.

Another common case of IPv6 fragmentation is the use of protocols like DNS and RADIUS, where the server response needs to be sent as a single UDP datagram. Network security policies MUST allow IPv6 fragments for permitted UDP traffic (e.g., DNS, RADIUS) where single-datagram responses are required. Allowing IPv6 fragments for permitted TCP traffic is RECOMMENDED unless the network infrastructure reliably performs TCP MSS clamping.

DISCUSSION: neither [RFC6145] nor [RFC6146] requires that NAT64 device performs MSS clamping, reducing MSS by 20 bytes while translating IPv6 to IPv4. Sounds like a useful feature though.

7.3.6. Representing IPv6 Addresses by CLAT

Certain CLAT implementations face challenges when translating incoming IPv6 packets with native (non-synthesized) source addresses (e.g. ICMPv6 packets sent by intermediate hops on the path). This lack of standardized translation mechanisms can lead to:

- Incomplete Traceroute: Omission of IPv6-only hops between the endpoint and NAT64 translator, hindering troubleshooting.

- Path MTU Discovery Issues: Potential disruptions in the PMTU discovery process.

DISCUSSION: shall the IETF consider standardizing a translation mechanism for such packets?

7.3.7. Custom DNS Configuration on Endpoints

In IPv6-mostly networks without PREF64 in RAs, hosts rely on DNS64 ([RFC7050]) to discover the NAT64 prefix for CLAT operation. Endpoints or applications configured with custom DNS resolvers (e.g., public or corporate DNS) may bypass the network-provided DNS64, preventing NAT64 prefix discovery and hindering CLAT functionality.

Where feasible, administrators SHOULD include PREF64 in RAs within IPv6-mostly networks to minimize reliance on DNS64. Administrators need to be aware of the potential for CLAT failures when endpoints use custom resolvers in environments lacking PREF64.

8. Security Considerations

9. Privacy Considerations

This document does not introduce any privacy considerations.

10. IANA Considerations

This memo does not introduce any requests to IANA.

11. References

11.1. 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/info/rfc2119>.

[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/info/rfc6146>.

[RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, DOI 10.17487/RFC6333, August 2011, <https://www.rfc-editor.org/info/rfc6333>.

[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/info/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/info/rfc7050>.

[RFC7335] Byrne, C., "IPv4 Service Continuity Prefix", RFC 7335, DOI 10.17487/RFC7335, August 2014, <https://www.rfc-editor.org/info/rfc7335>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

[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/info/rfc8174>.

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

[RFC8585] Palet Martinez, J., Liu, H. M.-H., and M. Kawashima, "Requirements for IPv6 Customer Edge Routers to Support IPv4-as-a-Service", RFC 8585, DOI 10.17487/RFC8585, May 2019, <https://www.rfc-editor.org/info/rfc8585>.

[RFC8925] Colitti, L., Linkova, J., Richardson, M., and T. Mrugalski, "IPv6-Only Preferred Option for DHCPv4", RFC 8925, DOI 10.17487/RFC8925, October 2020, <https://www.rfc-editor.org/info/rfc8925>.

11.2. Informative References

[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/info/rfc1918>.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <https://www.rfc-editor.org/info/rfc4861>.

[RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <https://www.rfc-editor.org/info/rfc4862>.

[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/info/rfc6052>.

[RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation Algorithm", RFC 6145, DOI 10.17487/RFC6145, April 2011, <https://www.rfc-editor.org/info/rfc6145>.

[RFC7225] Boucadair, M., "Discovering NAT64 IPv6 Prefixes Using the Port Control Protocol (PCP)", RFC 7225, DOI 10.17487/RFC7225, May 2014, <https://www.rfc-editor.org/info/rfc7225>.

[I-D.ietf-v6ops-dhcp-pd-per-device] Colitti, L., Linkova, J., and X. Ma, "Using DHCPv6-PD to Allocate Unique IPv6 Prefix per Client in Large Broadcast Networks", Work in Progress, Internet-Draft, draft-ietf-v6ops-dhcp-pd-per-device-07, 26 February 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-v6ops-dhcp-pd-per-device-07>.

[I-D.link-v6ops-claton] Linkova, J. and T. Jensen, "464 Customer-side Translator (CLAT): Node Recommendations", Work in Progress, Internet-Draft, draft-link-v6ops-claton-02, 28 February 2024, <https://datatracker.ietf.org/doc/html/draft-link-v6ops-claton-02>.

[I-D.link-v6ops-gulla] Linkova, J., "Using Subnet-Specific Link-Local Addresses to Improve SLAAC Robustness", Work in Progress, Internet-Draft, draft-link-v6ops-gulla-01, 25 February 2024, <https://datatracker.ietf.org/doc/html/draft-link-v6ops-gulla-01>.

Acknowledgements

TBA

Author's Address

Jen Linkova
Google
1 Darling Island Rd
Pyrmont NSW 2009
Australia

Email: furry13@gmail.com, furry@google.com