Workgroup: IPv6 Maintenance
Internet-Draft:
draft-collink-6man-pio-pflag-00
Updates: 4861, 4862 (if approved)
Published: 27 March 2023
Intended Status: Standards Track
Expires: 28 September 2023
Authors: L. Colitti J. Linkova
Google, LLC Google

Signalling DHCPv6 Prefix Delegation Availability to Hosts

Abstract

This document defines a ‘P’ flag in the Prefix Information Option of IPv6 Router Advertisements (RAs). The flag is used to indicate that the network prefers that hosts acquire global addresses using DHCPv6 PD instead of using SLAAC for this prefix.

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 28 September 2023.

Table of Contents

1. Introduction

IPv6 hosts, especially mobile hosts, usually have multiple global IPv6 addresses (e.g. stable addresses, privacy addresses, 464XLAT addresses, addresses for virtual systems etc).

On large networks, individually tracking these addresses can create scalability issues for the infrastructure, because routers must maintain multiple entries (neighbor cache, SAVI mappings, VXLAN routes, etc.) for each host. [I-D.collink-v6ops-ent64pd] discusses these challenges and proposes a solution that uses DHCPv6 PD [RFC8415].

On small networks, scaling to support multiple individual IPv6 addresses is less of a concern, because many home routers support hundreds of neighbor cache entries. On the other hand, address space is more limited compared to the number of hosts connected - the smallest home network might only have /60 prefixes, or even just a single /64.

A host cannot know in advance which address assignment method is most appropriate for the network, so there must be a mechanism for the network to communicate with this to the host.

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. Rationale

The information is passed to the host via a P flag in the Prefix Information Option (PIO). The reason is as follows:

* The information should be contained in the Router Advertisement because it must be available to the host before it decides to form IPv6 addresses from the prefix using SLAAC. Otherwise, the host might form IPv6 addresses from the PIO provided and start using them. This is suboptimal because if the host later acquires a prefix using DHCPv6 PD, it can either use both the prefix and SLAAC addresses, reducing the scalability benefits of using DHCPv6 PD, or can remove the SLAAC addresses, which would be disruptive for applications that are using them.

* This information is specific to the particular prefix being announced. For example, a network might want to assign global addresses via DHCPv6 PD, but use SLAAC for ULA addresses. Also, in a multihoming situation, one upstream network might choose to assign addresses via prefix delegation, and another via SLAAC.

4. Host Behaviour

4.1. Tracking and requesting prefixes

The host SHOULD NOT use SLAAC to obtain IPv6 addresses from prefix(es) with the P bit set.

For each network it is currently connected to, the host MUST keep a list of every PIO it has received with the P flag. Each time the client receives a Router Advertisement containing a PIO with the P bit set that is not in the list, and every time a previously-received PIO with the P bit set becomes deprecated:

* If the client has not previously received any delegated prefixes from the network, it SHOULD start DHCPv6 Prefix Delegation.

* If the client has already received delegated prefix(es) from one or more servers, it MUST send a RENEW request to each server, to obtain new prefixes. This allows the network to be renumbered.

Whenever a Prefix Information Option’s Valid lifetime reaches zero, or its P flag changes to 0, the prefix is removed from the list. When there are no such prefixes, the host SHOULD stop the DHCPv6 client if it has no other reason to run it. The lifetimes of any DHCPv6 prefixes already obtained are unaffected.

When a host requests a prefix via DHCPv6 PD, it MUST use the prefix length hint Section 18.2.4 of [RFC8415] to request a prefix that is short enough to form addresses via SLAAC. To ensure that all DHCP relays on link can act on the delegated prefix, the host SHOULD NOT use the Rapid Commit option.

The P flag is meaningless for link-local prefixes and any Prefix Information Option containing the link-local prefix MUST be ignored as specified in Section 5.5.3 of [RFC4862].

4.2. Using received prefix(es)

For every delegated prefix:

- The host MAY form as many IPv6 addresses from the prefix as it chooses.

- The host MAY use the prefix to provide IPv6 addresses to internal components such as virtual machines or containers.

- If the host is capable of acting as a router, and doing so is allowed by local policy, it MAY use the prefix to allow devices directly connected to it to obtain IPv6 addresses, e.g., by sending a Router Advertisement containing the prefix to a connected interface.

5. Multihoming

In multi-prefix multihoming, the host generally needs to associate the prefix with the router that advertised it (see for example, [RFC6724] Rule 5.5). If the host supports Rule 5.5, then it SHOULD associate each prefix with the link-local address of the DHCPv6 relay from which it received the packet.

6. Modifications to RFC-Mandated Behavior

6.1. Changes to RFC4861

This document makes the following changes to Section 4.6.2 of [RFC4861]

OLD TEXT:

==

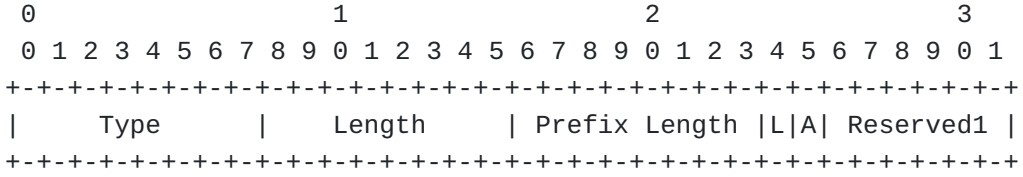

Figure 1

===

NEW TEXT

===

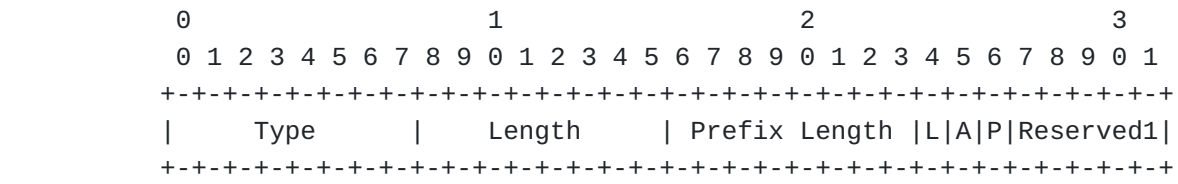

Figure 2

OLD TEXT

===

A 1-bit autonomous address-configuration flag. When set indicates that this prefix can be used for stateless address configuration as specified in [ADDRCONF].

Reserved1 6-bit unused field. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

===

NEW TEXT

===

A 1-bit autonomous address-configuration flag. When set indicates that this prefix can be used for stateless address configuration as specified in [ADDRCONF].

P 1-bit DHCPv6-PD flag. When set, indicates that this prefix SHOULD NOT be used for stateless address configuration. Instead the host SHOULD request a dedicated prefix via DHCPv6-PD and use that prefix for stateless address configuration.

Reserved1 5-bit unused field. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

===

6.2. Changes to RFC4862

This document makes the following changes to Section 5.5.3 of [RFC4862]:

OLD TEXT

===

For each Prefix-Information option in the Router Advertisement:

a) If the Autonomous flag is not set, silently ignore the Prefix Information option.

===

NEW TEXT

===

For each Prefix-Information option in the Router Advertisement:

a) If the P flag is set, start the DHCPv6 PD process and use the delegated prefix to assign addresses to the interfaces as described in draft-collink-6man-pio-pflag. The Prefix Information option SHOULD be processed as if A flag is set to zero.

b)If the Autonomous flag is not set, silently ignore the Prefix

===

7. Privacy Considerations

To be added

8. IANA Considerations

This memo includes no request to IANA.

9. Security Considerations

to be added

10. References

10.1. 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>.

[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>.

[RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, DOI 10.17487/RFC6724, September 2012, <https://www.rfc-editor.org/info/rfc6724>.

[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>.

[RFC8415] Mrugalski, T., Siodelski, M., Volz, B., Yourtchenko, A., Richardson, M., Jiang, S., Lemon, T., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 8415, DOI 10.17487/RFC8415, November 2018, <https://www.rfc-editor.org/info/rfc8415>.

10.2. Informative References

[I-D.collink-v6ops-ent64pd] Colitti, L., Linkova, J., and X. Ma, "Using DHCP-PD to Allocate /64 per Host in Broadcast Networks", Work in Progress, Internet-Draft, draft-collink-v6ops-ent64pd-02, 27 February 2023, <https://datatracker.ietf.org/doc/html/draft-collink-v6ops-ent64pd-02>.

Acknowledgements

Contributors

Authors' Addresses

Lorenzo Colitti
Google, LLC
Shibuya 3-21-3,
Japan

Email: lorenzo@google.com

Jen Linkova

Google
1 Darling Island Rd
Pyrmont NSW 2009
Australia

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