Internet Engineering Task Force (IETF)
Request for Comments: 9762
Updates: 4861, 4862
Category: Standards Track
ISSN: 2070-1721

L. Colitti
J. Linkova
X. Ma, Ed.
Google
D. Lamparter
NetDEF, Inc.

June 2025

Using Router Advertisements to Signal the Availability of DHCPv6 Prefix Delegation to Clients

Abstract

This document defines the P flag in the Prefix Information Option (PIO) of IPv6 Router Advertisements (RAs). The flag is used to indicate that the network prefers that clients use the deployment model in RFC 9663 instead of using individual addresses in the on-link prefix assigned using Stateless Address Autoconfiguration (SLAAC) or DHCPv6 address assignment.

This document updates RFC 4862 to indicate that the Autonomous flag in a PIO needs to be ignored if the PIO has the P flag set. It also updates RFC 4861 to specify that the P flag indicates DHCPv6 prefix delegation support for clients.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at https://www.rfc-editor.org/info/rfc9762.

Table of Contents

1. Introduction

[RFC9663] documents an IPv6 address assignment model where IPv6 devices obtain dedicated prefixes from the network via DHCPv6 prefix delegation (DHCPv6-PD) [RFC8415]. This model provides devices with a large IPv6 address space they can use to create addresses for communication, individually number virtual machines (VMs) or containers, or extend the network to downstream devices. It also provides scalability benefits on large networks because network infrastructure devices do not need to maintain per-address state, such as IPv6 neighbor cache, Source Address Validation Improvement (SAVI) [RFC7039] mappings, Virtual eXtensible Local Area Network (VXLAN) [RFC7348] routes, etc.

On networks with fewer devices, however, this model may not be appropriate, because scaling to support multiple individual IPv6 addresses per device is less of a concern. Also, many home networks currently offer prefix delegation but assume that a limited number of specialized devices and/or applications will require delegated prefixes and thus do not allocate enough address space to offer prefixes to every device that connects to the network. For example, if clients assume the [RFC9663] deployment model on a home network that only receives a /60 from the ISP and each client obtains a /64 prefix, then the network will run out of prefixes after 15 devices have been connected.

Therefore, to safely roll out the support of the deployment model defined in [RFC9663] on the client side, it is necessary to have a mechanism for the network to signal to the client which address assignment method is preferred.

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

Node: a device that implements IPv6 [RFC8200]

Host: any node that is not a router [RFC8200]

Client: a node that connects to a network and acquires addresses. The node may wish to obtain addresses for its own use, or it may be a router that wishes to extend the network to its physical or virtual subsystems, or both. It may be either a host or a router as defined by [RFC8200].

DHCPv6-PD: DHCPv6 Prefix Delegation [RFC8415]; a mechanism to delegate IPv6 prefixes to clients.

DHCPv6 IA_NA: Identity Association for Non-temporary Addresses (Section 21.4 of [RFC8415])

DHCPv6 IA_PD: Identity Association for Prefix Delegation (Section 21.21 of [RFC8415])

ND: Neighbor Discovery [RFC4861]

On-link address: an address that is assigned to an interface on a specified link [RFC4861]

On-link prefix: a prefix that is assigned to a specified link

Off-link: the opposite of "on-link" (see [RFC4861])

PIO: Prefix Information Option [RFC4862]

RA: Router Advertisement [RFC4861]

SLAAC: Stateless Address Autoconfiguration [RFC4862]

4. Rationale

The network administrator might want to indicate to clients that requesting a prefix via DHCPv6-PD and using that prefix for address assignment (see [RFC9663]) should be preferred over using individual addresses from the on-link prefix. The information is passed to the client via a P flag in the PIO. The reasons for it being a PIO flag are as follows:

* The information must be contained in the RA because it must be available to the client before it decides to form IPv6 addresses from the PIO prefix using SLAAC. Otherwise, the client might use SLAAC to form IPv6 addresses from the PIO provided and start using them, even if a unique per-client prefix is available via DHCPv6-PD. Forming addresses via SLAAC is suboptimal because if the client later acquires a prefix using DHCPv6-PD, it can either 1) use both the prefix and SLAAC addresses, reducing the scalability benefits of using DHCPv6-PD, or 2) 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 administrator might want clients to assign global addresses from delegated prefixes but form Unique Local IPv6 Unicast Addresses [RFC4193] from another PIO in the RA using SLAAC. Also, in a multihoming situation, one upstream network might choose to assign prefixes via prefix delegation and another via PIOs.

Note that setting the P flag in a PIO expresses the network operator’s preference that clients should attempt using DHCPv6-PD instead of performing individual address configuration on the prefix. For clients that honor this preference by requesting prefix delegation, the actual delegated prefix will necessarily be a prefix different from the one from the PIO.

5. P Flag Overview

The P flag (also called the DHCPv6-PD Preferred Flag) is a 1-bit PIO flag, located after the R flag [RFC6275]. The presence of a PIO with the P flag set indicates that the network prefers that clients use prefix delegation instead of acquiring individual addresses via SLAAC or DHCPv6 address assignment. This implies that the network has a DHCPv6 server capable of making DHCPv6 prefix delegations to every device on the network, as described in [RFC9663].

Figure 1 shows the resulting format of the PIO.

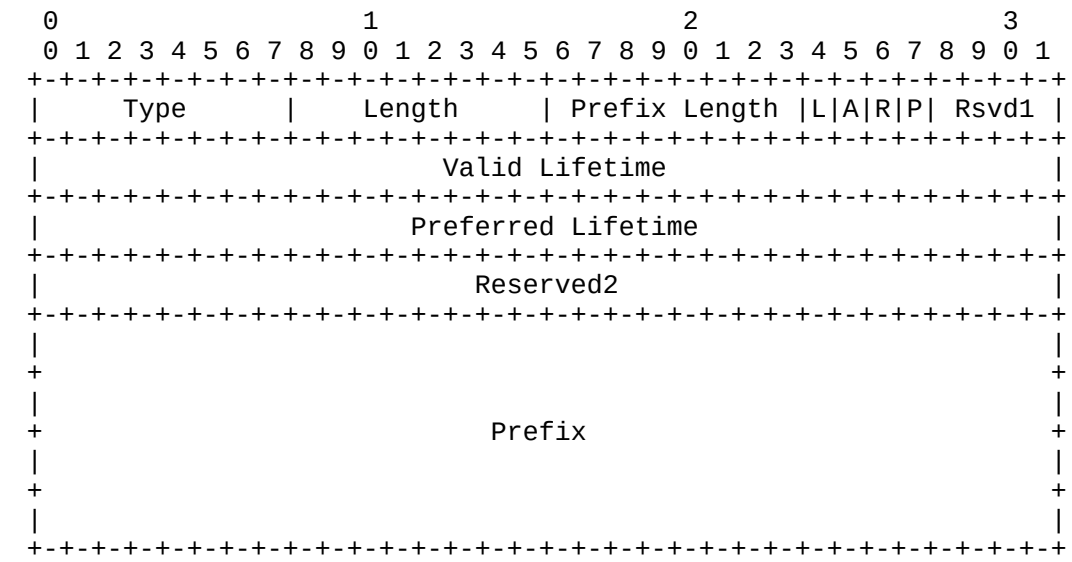

Figure 1

The P flag is independent of the value of the M and O flags in the RA. If the network desires to delegate prefixes to devices that support DHCPv6 prefix delegation but do not support the P flag, it SHOULD also set the M or O bits in the RA to 1, because some devices, such as Customer Edge (CE) routers [RFC7084], might not initiate DHCPv6 prefix delegation if both the M and O bits are set to zero.

6. Router Behavior

Routers SHOULD set the P flag to zero by default, unless explicitly configured by the administrator, and SHOULD allow the operator to set the P flag value for any given prefix advertised in a PIO. Routers MUST allow the P flag to be configured separately from the A flag. In particular, enabling or disabling the P flag MUST NOT trigger automatic changes in the A flag value set by the router.

7. Client Behavior

7.1. Processing the P Flag

This specification only applies to clients that support DHCPv6 prefix delegation. Clients that do not support DHCPv6 prefix delegation MUST ignore the P flag. The P flag is meaningless for link-local prefixes, and any PIO containing the link-local prefix MUST be ignored as specified in Section 5.5.3 of [RFC4862]. In the following text, all prefixes are assumed not to be link-local.

For each interface, the client MUST keep a list of every prefix that was received from a PIO with the P flag set and currently has a non-zero preferred lifetime. The list affects the behavior of the DHCPv6 client as follows:

* When a prefix's preferred lifetime becomes zero, either because the preferred lifetime expires or because the client receives a PIO for the prefix with a zero preferred lifetime, the prefix MUST be removed from the list.

* When the length of the list increases to one, the client SHOULD start requesting prefixes via DHCPv6 prefix delegation unless it is already doing so.

* When the length of the list decreases to zero, the client SHOULD stop requesting or renewing prefixes via DHCPv6 prefix delegation if it has no other reason to do so. The lifetimes of any prefixes already obtained via DHCPv6 are unaffected.

* If the client has already received delegated prefix(es) from one or more servers, then any time one or more prefix(es) are added to or removed from the list, the client MUST consider this to be a change in configuration information as described in Section 18.2.12 of [RFC8415]. In that case, the client MUST perform a REBIND, unless the list is now empty. This is in addition to performing a REBIND in the other cases required by that section. Issuing a REBIND allows the client to obtain new prefixes if necessary, for example, when the network is being renumbered. It also refreshes the state related to the delegated prefix(es).

When a client 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.

In order to achieve the scalability benefits of using DHCPv6-PD, the client SHOULD prefer to form addresses from the delegated prefix instead of using individual addresses in the on-link prefix(es). Therefore, when the client requests a prefix using DHCPv6-PD, the client SHOULD NOT use SLAAC to obtain IPv6 addresses from PIOs with the P and A bits set. Similarly, if all PIOs processed by the client have the P bit set, the client SHOULD NOT request individual IPv6 addresses from DHCPv6, i.e., it SHOULD NOT include any IA_NA options in Solicit messages [RFC8415]. The client MAY continue to use addresses that are already configured.

If the client does not obtain any suitable prefixes via DHCPv6-PD that are suitable for SLAAC, it MAY choose to disable further processing of the P flag on that interface, allowing the client to fall back to other address assignment mechanisms, such as forming addresses via SLAAC (if the PIO has the A flag set to 1) and/or requesting individual addresses via DHCPv6.

7.2. Using Delegated Prefix(es)

If the delegated prefix is too long to be used for SLAAC, the client MUST ignore it, as Section 7 of [RFC9663] requires the network to provide a SLAAC-suitable prefix to clients. If the prefix is shorter than required for SLAAC, the client SHOULD accept it, allocate one or more longer prefixes suitable for SLAAC, and use the prefixes as described below.

For every accepted prefix:

* The client MAY form as many IPv6 addresses from the prefix as it chooses.

* The client MAY use the prefix to provide IPv6 addresses to internal components such as VMs or containers.

* The client MAY use the prefix to allow devices directly connected

to it to obtain IPv6 addresses. For example, the client MAY route traffic for that prefix to an interface and send an RA containing a PIO for the prefix on that interface. That interface MUST NOT be the interface the prefix is obtained from. If the client advertises the prefix on an interface and it has formed addresses from the prefix, then it MUST act as though the addresses were assigned to that interface for the purposes of Neighbor Discovery and Duplicate Address Detection.

The client MUST NOT send or forward packets with destination addresses within a delegated prefix to the interface that it obtained the prefix on, as this can cause a routing loop. This problem will not occur if the client has assigned the prefix to another interface. Another way the client can prevent this problem is to add to its routing table a high-metric discard route for the delegated prefix.

7.3. Absence of PIOs with the P Bit Set

The P bit is purely a positive indicator, telling nodes that DHCPv6 prefix delegation is available and the network prefers that nodes use it, even if they do not have any other reason to run a prefix delegation client. The absence of any PIOs with the P bit does not carry any kind of signal to the opposite and MUST NOT be processed to mean that DHCPv6-PD is absent. In particular, nodes that run DHCPv6-PD due to explicit configuration or by default (e.g., to extend the network) MUST NOT disable DHCPv6-PD on the absence of PIOs with the P bit set. A very common example of this are CE routers as described by [RFC7084].

7.4. On-Link Communication

When the network delegates unique prefixes to clients, each client will consider other clients' destination addresses to be off-link, because those addresses are from the delegated prefixes and are not within any on-link prefix. When a client sends traffic to another client, packets will initially be sent to the default router. The router may respond with an ICMPv6 redirect message (Section 4.5 of [RFC4861]). If the client receives and accepts the redirect, then traffic can flow directly from device to device. Therefore, hosts supporting the P flag SHOULD process redirects unless configured otherwise. Hosts that do not process ICMPv6 redirects, and routers that do not act on ICMPv6 redirects, may experience higher latency while communicating to prefixes delegated to other clients on the same link.

7.5. Source Address Selection

For the purpose of source address selection [RFC6724], if the host creates any addresses from a delegated prefix, it SHOULD treat those addresses as if they were assigned to the interface on which the prefix was received. This includes placing them in the candidate set and associating them with the outgoing interface when implementing Rule 5 of the source address selection algorithm [RFC6724].

8. Multihoming

In multi-prefix multihoming, the host generally needs to associate the prefix with the router that advertised it (for example, see Rule 5.5 in [RFC6724]). If the host supports Rule 5.5, then it SHOULD associate each prefix with the link-local address of the DHCPv6 server or relay from which it received the REPLY packet. When receiving multiple REPLYs carrying the same prefix from distinct link-local addresses, the host SHOULD associate that prefix with all of these addresses. This can commonly happen in networks with redundant routers and DHCPv6 servers or relays.

9. Modifications to RFC-Mandated Behavior

9.1. Changes to RFC 4861

This document makes the following changes to Section 4.2 of [RFC4861]:

OLD TEXT:

> Note: If neither M nor O flags are set, this indicates that no information is available via DHCPv6.

NEW TEXT:

> Note: If the M, O, or P (RFC 9762) flags are not set, this indicates that no information is available via DHCPv6.

9.2. Changes to RFC 4862

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.
>
> b) If the prefix is the link-local prefix, silently ignore the Prefix Information option.
>
> c) If the preferred lifetime is greater than the valid lifetime, silently ignore the Prefix Information option. A node MAY wish to log a system management error in this case.
>
> d) If the prefix advertised is not equal to the prefix of an address configured by stateless autoconfiguration already in the list of addresses associated with the interface (where "equal" means the two prefix lengths are the same and the first prefix-length bits of the prefixes are identical), and if the Valid Lifetime is not 0, form an address (and add it to the list) by combining the advertised prefix with an interface identifier of the link as follows:

NEW TEXT:

> For each Prefix Information Option in the Router Advertisement:
>
> a) If the prefix is the link-local prefix, silently ignore the Prefix Information Option.
>
> b) If the P flag is set and the node implements RFC 9762, it SHOULD treat the Autonomous flag as if it was unset and use prefix delegation to obtain addresses as described in RFC 9762.
>
> c) If the Autonomous flag is not set, silently ignore the Prefix Information Option.
>
> d) If the preferred lifetime is greater than the valid lifetime, silently ignore the Prefix Information Option. A node MAY wish to log a system management error in this case.
>
> e) If the prefix advertised is not equal to the prefix of an address configured by stateless autoconfiguration already in

| the list of addresses associated with the interface (where
| "equal" means the two prefix lengths are the same and the
| first prefix-length bits of the prefixes are identical) and if
| the Valid Lifetime is not 0, form an address (and add it to
| the list) by combining the advertised prefix with an interface
| identifier of the link as follows:

10. Security Considerations

The mechanism described in this document relies on the information provided in the RA and therefore shares the same security model as SLAAC. If the network does not implement RA-Guard [RFC6105], an attacker might send RAs containing the PIO used by the network, set the P flag to 1, and force hosts to ignore the A flag. In the absence of DHCPv6-PD infrastructure, hosts would either obtain no IPv6 addresses or, if they fall back to other IPv6 address assignment mechanisms such as SLAAC and IA_NA, would experience delays in obtaining IPv6 addresses. If the network does not support DHCPv6-Shield [RFC7610], the attacker could also run a rogue DHCPv6 server, providing the host with invalid prefixes or other invalid configuration information.

The attacker might force hosts to oscillate between DHCPv6-PD and PIO-based SLAAC by sending the same set of PIOs with and then without the P flag set. That would cause the clients to issue REBIND requests, increasing the load on the DHCP infrastructure. However, Section 14.1 of [RFC8415] requires that DHCPv6-PD clients rate-limit transmitted DHCPv6 messages.

It should be noted that if the network allows rogue RAs to be sent, the attacker would be able to disrupt hosts' connectivity anyway, so this document doesn't introduce any fundamentally new security considerations.

Security considerations inherent to the PD-per-device model are documented in Section 15 of [RFC9663].

11. Privacy Considerations

The privacy implications of implementing the P flag and using DHCPv6-PD to assign prefixes to hosts are similar to the privacy implications of using DHCPv6 to assign individual addresses. If the DHCPv6 infrastructure assigns the same prefix to the same client, then an observer might be able to identify clients based on the highest 64 bits of the client's address. Those implications and recommended countermeasures are discussed in Section 13 of [RFC9663].

Implementing the P flag support on a host and receiving will enable DHCPv6 on that host if the host receives an RA containing a PIO with the P bit set. Sending DHCPv6 packets may reveal some minor additional information about the host, most prominently the hostname. This is not a new concern and would apply for any network that uses DHCPv6 and sets the M flag in RAs.

No privacy considerations result from supporting the P flag on the sender side.

12. IANA Considerations

IANA has made the following allocation in the "IPv6 Neighbor Discovery Prefix Information Option Flags" registry [RFC8425]:

PIO Option Bit	Description	Reference
3	P - DHCPv6-PD Preferred Flag	RFC 9762

+---------------+------------------------------+-----------+

Table 1

13. References

13.1. 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>.

[RFC8425] Troan, O., "IANA Considerations for IPv6 Neighbor Discovery Prefix Information Option Flags", RFC 8425, DOI 10.17487/RFC8425, July 2018, <https://www.rfc-editor.org/info/rfc8425>.

13.2. Informative References

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, DOI 10.17487/RFC4193, October 2005, <https://www.rfc-editor.org/info/rfc4193>.

[RFC6105] Levy-Abegnoli, E., Van de Velde, G., Popoviciu, C., and J. Mohacsi, "IPv6 Router Advertisement Guard", RFC 6105, DOI 10.17487/RFC6105, February 2011, <https://www.rfc-editor.org/info/rfc6105>.

[RFC6275] Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, DOI 10.17487/RFC6275, July 2011, <https://www.rfc-editor.org/info/rfc6275>.

[RFC7039] Wu, J., Bi, J., Bagnulo, M., Baker, F., and C. Vogt, Ed., "Source Address Validation Improvement (SAVI) Framework", RFC 7039, DOI 10.17487/RFC7039, October 2013, <https://www.rfc-editor.org/info/rfc7039>.

[RFC7084] Singh, H., Beebee, W., Donley, C., and B. Stark, "Basic Requirements for IPv6 Customer Edge Routers", RFC 7084, DOI 10.17487/RFC7084, November 2013,

<https://www.rfc-editor.org/info/rfc7084>.

[RFC7348] Mahalingam, M., Dutt, D., Duda, K., Agarwal, P., Kreeger, L., Sridhar, T., Bursell, M., and C. Wright, "Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", RFC 7348, DOI 10.17487/RFC7348, August 2014, <https://www.rfc-editor.org/info/rfc7348>.

[RFC7610] Gont, F., Liu, W., and G. Van de Velde, "DHCPv6-Shield: Protecting against Rogue DHCPv6 Servers", BCP 199, RFC 7610, DOI 10.17487/RFC7610, August 2015, <https://www.rfc-editor.org/info/rfc7610>.

[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>.

[RFC9663] Colitti, L., Linkova, J., Ed., and X. Ma, Ed., "Using DHCPv6 Prefix Delegation (DHCPv6-PD) to Allocate Unique IPv6 Prefixes per Client in Large Broadcast Networks", RFC 9663, DOI 10.17487/RFC9663, October 2024, <https://www.rfc-editor.org/info/rfc9663>.

Acknowledgements

Thanks to Nick Buraglio, Brian Carpenter, Tim Chown, David Farmer, Fernando Gont, Susan Hares, Mahesh Jethanandani, Suresh Krishnan, Ted Lemon, Andrew McGregor, Tomek Mrugalski, Erik Nordmark, Michael Richardson, Patrick Rohr, John Scudder, Ole Tran, Dirk Von Hugo, ric Vyncke and Timothy Winters for the discussions, reviews, input, and contributions.

Authors' Addresses

Lorenzo Colitti
Google
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

Xiao Ma (editor)
Google
Shibuya 3-21-3,
Japan
Email: xiaom@google.com

David 'equinox' Lamparter
NetDEF, Inc.
04229 Leipzig
Germany
Email: equinox@diac24.net, equinox@opensourcerouting.org