

Network Working Group
Internet-Draft
Intended status: Informational
Expires: May 10, 2019

Naveen Kottapalli, Ed.
Benu Networks
November 6, 2018

IPv6 Stateless Prefix Management
draft-naveen-slaac-prefix-management-00

Abstract

The Dynamic Host Configuration Protocol for IPv6 (DHCPv6) provides a provision for the hosts to request for a specific IPv6 address and manage the same, whereas the Stateless Address AutoConfiguration (SLAAC) doesn't have such a provision. Also, unavailability of DHCPv6 across all OS platforms has limited the IPv6 nodes to not use features like soliciting a prefix of desired length and lifetime, renewing, releasing / declining a prefix, etc. This document proposes IPv6ND extensions for enabling SLAAC devices to solicit prefix information as a hint PIO in RS and other methods like renewing or releasing or declining a prefix without introducing any new ICMPv6 message or option types.

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 May 10, 2019.

Copyright Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of

publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Terminology	3
2.1.	General	3
2.1.1.	/64 Prefix	3
2.1.2.	Non /64 Prefix	3
2.1.3.	Requesting Node	3
2.1.4.	Delegating Node or Router	4
2.1.5.	Requirements Language	4
3.	Related works	4
4.	Use cases	4
4.1.	Host Soliciting a Prefix	4
4.2.	CPE or Downstream Router Soliciting a Prefix	5
5.	Soliciting Prefix	6
5.1.	Requesting Node Specification	6
5.2.	Delegating Router Specification	8
6.	Renewing Prefix	9
6.1.	Requesting Node Specification	10
6.2.	Delegating Router Specification	10
7.	Releasing Prefix	11
7.1.	Requesting Node Specification	11
7.2.	Delegating Router Specification	11
8.	Declining Prefix	12
8.1.	Requesting Node Specification	12
8.2.	Delegating Router Specification	12
9.	Processing Router Advertisements	12
10.	Unsolicited Router Advertisements	12
10.1.	Requesting Node Specification	12
10.2.	Delegating Router Specification	12
11.	Backward Compatibility	13
12.	IANA Considerations	13
13.	Security Considerations	13
14.	Acknowledgements	13
15.	References	14
15.1.	Normative References	14
15.2.	Informative References	14
Appendix A.	Comparison with other works	15
Author's Address		15

1. Introduction

IPv6 Stateless Address AutoConfiguration (SLAAC) [ADDRCONF] in its current form doesn't offer services like soliciting a prefix of specific length or for a stipulated time forces the nodes to be dictated solely by the routers that assigns the prefixes. Though a SLAAC node uses the prefix received from a router, at times a host might want to make sure that the prefix is still valid with the router and may want to renew the prefix with router. Similarly, since every prefix that the router assigns to a node is a resource that costs the router w.r.t routing and subscriber management, the router also wants to release the prefixes back to free pool those are not being used by the SLAAC nodes.

This document presents an IPv6 ND based messaging protocol as an extension for SLAAC nodes that do similar things as does DHCPv6 protocol. For example, soliciting or renew or release or decline a prefix using existing standards.

2. Terminology

2.1. General

This document uses the same terminology defined in IPv6 Neighbor Discovery [RFC 4861](#) [RFC4861], IPv6 Stateless Address Autoconfiguration [RFC 4862](#) [RFC4862]

2.1.1. /64 Prefix

An IPv6 prefix of length 64 bits, which will be used by the end nodes for preparing and using a /128 bit address.

2.1.2. Non /64 Prefix

An IPv6 prefix of length less than 64 bits, which will be used by nodes as a pool of prefixes for allocating either /64 prefixes or a non /64 prefix (subnet of the prefix pool) to the downstream nodes.

2.1.3. Requesting Node

A requesting node is a node that can be either a host / CPE / a downstream router etc. that is soliciting a prefix from router. If a host (end device) is soliciting a prefix, the prefix length will always be /64-bit, whereas if a CPE / downstream router is soliciting a prefix it can be a non /64-bit.

2.1.4. Delegating Node or Router

A delegating node or router is a node that assigns either 64-bit length prefixes or non 64-bit length prefix to requesting nodes. A delegating node or router MAY also do subscriber management along with the routing functionality.

2.1.5. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

3. Related works

A few drafts regarding prefix negotiation, prefix return were defined as a protocol set earlier.

In [[NDPD](#)], authors proposed a new PI option with a list of operation types that can be exchanged with the router.

In [[HPD](#)], authors proposed new ICMPv6 message types that can be exchanged with the router for prefix negotiation.

In [[PIO-X](#)], authors proposed new flag 'X' to be set in PIO flags field, which indicates that prefix received in a PIO is for host's exclusive use.

In [[UNIFIED](#)], the author proposes including a PIO option in RS messages. This document uses the same mechanism, but without including any new PIO flag bits.

The mechanism described in this document though follows some similar approach of other works, no new ICMPv6 messages or new option types are introduced, and not even new flags to be set in the PIO option. This specification uses the existing messages types and PIO option to achieve the objective.

4. Use cases

4.1. Host Soliciting a Prefix

A host SHALL solicit one or more 64-bit length prefixes [[RFC7934](#)] from a router; following can be the cases where a host wants to solicit prefixes from the available routers.

- o A host wants to use any prefix given by router only for a stipulated period.

- o A host wants to use a prefix only for a stipulated time period.
- o A host wants to use the same prefix that was received earlier from router. An example of this can be like when mobility happens the host solicits for the same prefix that was used earlier.
- o A host wants to assign a specific on-link prefix to the interface (the way how a host gets the prefix information is out of scope of this document).
- o A host wants to assign a specific off-link prefix to the interface (the way how a host gets the prefix information is out of scope of this document).
- o A host wants to renew an existing prefix with the same router from where the prefix was received.




Figure 1

4.2. CPE or Downstream Router Soliciting a Prefix

A CPE or downstream router SHALL solicit one or more non /64-bit length prefixes [[RFC7934](#)] from routers and assigns either /64-bit or a subnet from the received prefix to downstream nodes. Following can be the cases where a CPE or downstream router wants to solicit prefixes from the available routers.

- o A CPE or downstream router wants to use any non /64-bit prefix for a stipulated time period.
- o A CPE or downstream router wants to use a non /64-bit prefix for a stipulated time period.
- o A CPE or downstream router wants to use a specific on-link non /64-bit prefix of length 'n' bits, where 'n' is less than 64 bits.
- o A CPE or downstream router wants to use a specific off-link non /64-bit prefix of length 'n' bits, where 'n' is less than 64 bits.
- o A CPE or downstream router wants to renew an existing prefix with the same router from where the prefix was received.




Figure 2

A CPE or downstream router MAY decide on the length of non /64 prefix based on the number of downstream nodes that a CPE or downstream router has to cater the services. If the CPE or downstream router has a requirement to cater 256 downstream nodes, then the prefix length inside a PIO will be set to 56. The logic of how a CPE or downstream router knows the prefix length is out of scope of this document.

5. Soliciting Prefix

5.1. Requesting Node Specification

Requesting nodes that are compliant with this specification MAY append PIO option(s) in RS [UNIFIED], which can be triggered when one of the events mentioned in IPv6 Neighbor Discovery [RFC4861] occurs.

A requesting node MAY set the desired length of prefix it wants to solicit in the length field of each PIO that is included in RS. In the absence of a valid length field i.e., when the prefix length of a PIO is set to 0, by default the router assumes the length to be a 64-bit.

A requesting node MAY set the following flags of the PIO option when soliciting a prefix.

L - 1-bit on-link flag. When set to 1, it means that a requesting node is soliciting a prefix that is considered as on-link. When

set to 0, it means that a requesting node is soliciting a prefix that can be considered as off-link.

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

R - TODO - To be filled with use case.

A requesting node MUST set PIO's valid lifetime field to a positive integer value that is intended to be used or 0xffffffff. While soliciting a prefix a requesting node MUST NOT set the valid lifetime field of a PIO with 0.

A requesting node MAY set multiple desired prefixes that are intended to be used in each of the PIO option included in RS. In the absence of valid prefix in a PIO i.e. when the requesting node do not have any prefix information to be set in a PIO, value of prefix MUST be set to '::', while the other fields of PIO like prefix length, flags and valid lifetime are set. Since a NULL prefix in a PIO ('::') indicates a router that the requesting node is ready to take any prefix(es) that router assigns, a requesting node MUST NOT include more than one NULL prefix PIOs('::') in a RS.

An example of when a requesting node might include prefix in a PIO is, a node already has a prefix and due to mobility event the node triggers another RS. In such a case since the node already has a valid prefix which is under use, the same prefix SHALL be added in a PIO of RS that is being sent towards router.

An example of when a requesting node does not include prefix in a PIO is, a requesting node's interface is brought up without any prior prefix information available and the requesting node is looking for any of either /64 bit prefix or non /64 bit prefix for a stipulated time.

If the requesting node already has the information about the router from where the prefix was received earlier, RS with PIO included MAY be unicast to that particular router. A requesting node SHOULD transmit up to MAX_RTR_SOLICITATIONS [RFC4861] RS messages, each separated by at least RTR_SOLICITATION_INTERVAL [RFC4861] seconds.

In the absence of RA even after maximum attempts or on receipt of RA with the prefixes that requesting node has solicited with 0 life time MUST be treated as the solicited prefixes (if any) are no more available for the requesting node and all such prefixes MUST be discarded. Also, it is highly discouraged to not send another RS with prefixes in PIOs for which there is no RA or RA is received with

life time as 0. If a requesting node does not get RA for all the PIOs even after MAX_RTR_SOLICITATIONS times, the requesting node MUST send a RS by excluding the PIO option(s).

Requesting nodes while soliciting prefix(es), SHALL use the Nonce option [SEND] in RS message as transaction id to track the request and response messages.

5.2. Delegating Router Specification

The routers that are compliant with this specification look for the PIO option(s) present in RS and checks if a prefix can be assigned to the requesting node or not.

If prefix length in a PIO option is set to non-zero value, router checks its configuration if the desired prefix length can be allocated or not. If the prefix length of a PIO is set to a value less than 64 and if the router configuration doesn't allow prefix allocation of non /64 or the desired prefix length, then the PIO option MUST be ignored. If prefix length in a PIO option is set to zero (0), then router assumes that the requesting node is soliciting a /64 prefix.

A router checks for the following flags in the PIO option when present in RS.

L - When set to 1, the router checks if the solicited prefix can be assigned to the device or not which considers the prefix as "on-link". If the router's configuration doesn't allow "on-link" prefix allocation, then the PIO option MUST be ignored. When set to 0, the router checks if the solicited prefix can be assigned to the device or not which considers the prefix as "off-link". If the router's configuration doesn't allow "off-link" prefix allocation, then the PIO option MUST be ignored.

A - The router checks if the solicited prefix can be assigned to the device or not which uses the prefix for autonomous auto address configuration. If the router's configuration doesn't allow prefix allocation for autonomous auto address configuration, then the PIO option MUST be ignored.

R - TODO - To be filled with use case.

If valid lifetime inside a PIO is set to 0xffffffff, then routers configured lifetime will be considered. If valid lifetime inside PIO is set to a positive integer value that a requesting node wants to use, router checks its own configuration and if the solicited lifetime is less than or equal to the router's allowed configuration,

then the solicited lifetime SHALL be granted to the PIO and will be included in RA. If the router's configured lifetime is less than solicited lifetime, and there is no barring of such PIOs, router SHALL include the router's configured lifetime in RA. If the valid lifetime of a PIO that is chosen based on above logic is less than the router's preferred lifetime configuration, then the valid lifetime value itself SHALL be copied to preferred lifetime of PIO.

If the prefix inside a PIO option is non NULL i.e. some valid prefix is received, then router checks if the solicited prefix with the requested length is available or not. If such prefix is available and can be allocated, then the router SHALL assign the prefix and include in the PIO option of RA. If the solicited prefix is not available, then PIO option MUST be ignored. If the prefix inside a PIO option is a NULL prefix ('::'), then router SHALL assign one of the prefixes that are available and include the PIO option in RA. If there are more than one NULL prefix PIOs present in RS, then the router MUST process only one NULL prefix PIO ('::') and rest of NULL prefix PIOs SHALL be ignored.

If a router can assign any one of the solicited prefixes with the requested prefix length and lifetime, a RA SHALL be sent back to the requesting node with such PIOs included. For any reason, if a router is not able to assign any of the prefixes from the PIO option(s) received in RS, no RA SHALL be sent by router. This lets the other routers who can assign the prefix to respond and the network will not be choked with unnecessary RAs.

A router MUST send a deprecated PIO in RA to the received PIO of RS in the following cases.

- o RS is unicast to the router and the prefix inside a PIO is non NULL
- o
 - * PIO cannot be honored because of either prefix length or lifetime or flags
 - * PIO cannot be honored because of unavailability of solicited prefix

6. Renewing Prefix

6.1. Requesting Node Specification

Requesting nodes that are compliant with this specification MAY renew a prefix based on events such as:

- o Prefix is used for 1/3rd of the prefix lifetime
- o Any other period that a requesting node wants to renew at

Requesting nodes that are compliant with this specification when trying to renew a prefix, shall prepare the RS message the same way as mentioned in the Soliciting section and the RS message is unicast to the router from where the prefixes were received. A requesting node SHOULD transmit up to MAX_RTR_SOLICITATIONS [[RFC4861](#)] RS messages, each separated by at least RTR_SOLICITATION_INTERVAL [[RFC4861](#)] seconds. In the absence of RA even after maximum attempts, a requesting node MAY continue using the prefix till the prefix lifetime expires. In case when a RA is received with zero lifetime for the prefixes that requesting node had solicited; those prefixes MUST be treated as no more available for the requesting node and all such prefixes MUST be discarded immediately. Also, it is highly discouraged to not send the prefixes in PIOs for which there is no RA or RA is received with life time as 0. If a requesting node does not get RA for all the PIOs even after MAX_RTR_SOLICITATIONS times, the requesting node MUST send a RS by excluding the PIO option(s).

Requesting nodes while renewing prefix(es), SHALL use the Nonce option [SEND] in RS message as transaction id to track the request and response messages.

6.2. Delegating Router Specification

When a router receives a RS message that is unicast and has valid prefix information in PIOs, router checks if the prefix lease context of requesting node is already present in its lease database or not. If the lease context is present; it means that the requesting node is trying to renew the prefix. If the lease context is not present, the RS MUST be treated as solicitation and process the RS as described in Solicitation section. While renewing a prefix of the requesting node routers SHALL just refer the prefix lifetime and ignore the other fields of PIO that were considered during solicitation.

If prefix valid lifetime inside a PIO is set to a non-zero value that a requesting node wants to use, router checks its own configuration and if the solicited lifetime + elapsed lifetime is less than or equal to the routers total allowed configuration, then the solicited lifetime SHALL be granted to the PIO and will be included in RA. If the router's configured lifetime is less than solicited lifetime +

elapsed lifetime and if the router's configuration restricts barring such PIOs, the PIO option SHALL be ignored.

If the valid lifetime of a PIO that is chosen based on above logic is less than the router's preferred lifetime configuration, then the valid lifetime value itself SHALL be copied to preferred lifetime of PIO.

7. Releasing Prefix

7.1. Requesting Node Specification

Requesting nodes that are compliant with this specification MAY release a prefix based on events such as:

- o Lifetime of the prefix got expired and the requesting node does not want to use the prefix anymore.
- o Interface to which the prefix got assigned is brought down

A requesting node that wants to release prefixes back to router MUST set the lifetime of each prefix to zero (0) in PIO option of RS and MUST unicast to the router from where the prefixes were received.

Requesting nodes while releasing prefix(es), SHALL use the Nonce option [SEND] in RS message as transaction id to track the request and response messages.

7.2. Delegating Router Specification

When a router receives a RS message that is unicast and has valid prefix information in PIOs with valid lifetime as zero (0), router checks if the prefix lease context of requesting node is already present in its lease database or not. If the lease context is present, it means that the requesting node is trying to release the prefix. If the lease context is not present, then the PIO SHALL be ignored. Once the prefix context of the requesting node is released, the prefix SHALL be marked as free and be made available for further prefix allocations.

If a router is able to release a prefix with the mentioned prefix length, a RA SHALL be sent back to the node with the same PIO option and the lifetime set to 0. If a router is not able to release any of the prefixes from the PIO option(s) received in RS, no RA SHALL be sent by router.

8. Declining Prefix

8.1. Requesting Node Specification

Requesting nodes that are compliant with this specification MAY decline a prefix based on events such as:

- o Solicited prefix length is not assigned by a router.
- o Solicited lifetime of the prefix is not assigned by a router.

Above list of events is not an exhaustive list of cases where the requesting nodes decline the prefix and can be for any other reasons as well. Requesting nodes follow the same method described in Releasing section to decline a prefix.

8.2. Delegating Router Specification

The routers that are compliant with this specification treats a decline as release and follow the same method described in Releasing section.

9. Processing Router Advertisements

When a requesting node receives a RA from a router with PIO included, the router's address SHALL be stored against each prefix and process the RA using the procedure mentioned in [[RFC4861](#)]. The stored router's address can be used while renewing or releasing or declining the prefix.

10. Unsolicited Router Advertisements

10.1. Requesting Node Specification

When a requesting node receives an unsolicited RA from router for which no RS was sent, that RA must be treated as an unsolicited RA and SHALL be processed the same way defined in above sections. If a requesting node doesn't like the prefix lifetime received in unsolicited RA, a RS SHALL be triggered to renew the prefix.

10.2. Delegating Router Specification

Routers that are compliant with this specification follow the same method defined in [[RFC4861](#)] for sending unsolicited router advertisements.

11. Backward Compatibility

Routers that are compliant with this specification SHALL have the below configuration elements to decide the router behaviour.

- o Honoring the PIOs in RS - When a router is configured to not honor the PIOs in RS, all the PIOs SHALL be ignored and the RS MUST be processed as defined in [\[RFC4861\]](#).
- o Honoring non /64 prefix length PIOs in RS - When a router is configured to not honor the PIOs that request non /64 prefix length in RS, such PIOs SHALL be ignored and the RS MUST be processed as defined in [\[RFC4861\]](#).
- o Supported length of non /64 prefixes - A router can be configured with the non /64 prefix lengths that are supported, so that the PIOs in RS message can be honored.
- o Ignore PIOs with lifetime greater than configured value - A router can be configured to not honor the PIOs that solicit the prefix lifetime more than the configured router's prefix lifetime. When this flag is disabled, if a PIO in RS requests for a lifetime greater than the configured value, then the router's configured value shall be included in PIO of RA.

12. IANA Considerations

This memo includes no request to IANA.

13. Security Considerations

Security considerations for IPv6 Neighbor Discovery [\[RFC4861\]](#) and DHCPv6 [\[RFC3315\]](#)[\[RFC3633\]](#) apply to this document.

SEcure Neighbor Discovery (SEND) [\[SEND\]](#) can provide authentication for RS/RA exchanges with no need for additional securing mechanisms.

14. Acknowledgements

This work was motivated by discussions on the 6man and v6ops list.

The following individuals contributed to the document: Fred Templin

The following individuals provided useful comments that improved the document: Alexandre Petrescu

15. References

15.1. 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>>.

15.2. Informative References

- [HPD] Byung-Yeob Kim, Kyeong-Jin Lee, Jung-Soo Park, Hyoung-Jun Kim, "Hierarchical Prefix Delegation Protocol for Internet Protocol Version 6 (IPv6)", February 2004, <<https://tools.ietf.org/html/draft-bykim-ipv6-hpd-01>>.
- [NDPD] A. Kaiser, S. Decremps, N. Oualha, A. Petrescu, "Prefix Delegation extension to Neighbor Discovery protocol", February 2013, <<https://tools.ietf.org/html/draft-kaiser-nd-pd-01>>.
- [PIO-X] Kline, E. and M. Abrahamsson, "IPv6 Router Advertisement Prefix Information Option eXclusive Flag", March 2017, <<https://www.ietf.org/internet-drafts/draft-pioxfolks-6man-pio-exclusive-bit-02.txt>>.
- [RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", [RFC 2629](#), DOI 10.17487/RFC2629, June 1999, <<https://www.rfc-editor.org/info/rfc2629>>.
- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3315](#), DOI 10.17487/RFC3315, July 2003, <<https://www.rfc-editor.org/info/rfc3315>>.
- [RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", [BCP 72](#), [RFC 3552](#), DOI 10.17487/RFC3552, July 2003, <<https://www.rfc-editor.org/info/rfc3552>>.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", [RFC 3633](#), DOI 10.17487/RFC3633, December 2003, <<https://www.rfc-editor.org/info/rfc3633>>.

- [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>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [RFC 5226](#), DOI 10.17487/RFC5226, May 2008, <<https://www.rfc-editor.org/info/rfc5226>>.
- [RFC7934] Colitti, L., Cerf, V., Cheshire, S., and D. Schinazi, "Host Address Availability Recommendations", [BCP 204](#), [RFC 7934](#), DOI 10.17487/RFC7934, July 2016, <<https://www.rfc-editor.org/info/rfc7934>>.
- [UNIFIED] F. Templin, "A Unified Stateful/Stateless Autoconfiguration Service for IPv6", September 2018, <<https://datatracker.ietf.org/doc/draft-templin-6man-dhcpv6-ndopt>>.

[Appendix A](#). Comparison with other works

TODO: Add the comparison details here.

Author's Address

Naveen Kottapalli (editor)
Benu Networks
300 Concord Road
Billerica, MA 01821
USA

Phone: +1 978 223 4700
Email: nkottapalli@benu.net

