IPv6 Maintenance (6man) Working Group
Internet-Draft
Updates: 4191, 4861, 4862, 8106, 9096 (if approved)
Intended status: Standards Track
Expires: 23 April 2026

F. Gont
SI6 Networks
J. Zorz
6connect
R. Patterson
Sky UK
J. Linkova, Ed.
Google

20 October 2025

Improving the Robustness of Stateless Address Autoconfiguration (SLAAC) to Flash Renumbering Events
draft-ietf-6man-slaac-renum-11

Abstract

In scenarios where network configuration information becomes invalid without explicit notification to the local network, local hosts may end up employing stale information for an unacceptably long period of time, thus resulting in interoperability problems. This document improves the reaction of IPv6 Stateless Address Autoconfiguration to such configuration changes. It formally updates RFC 4191, RFC 4861, RFC 4862, RFC 8106, and RFC 9096.

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 23 April 2026.

Copyright Notice

Table of Contents

1. Introduction

IPv6 Stateless Address Autoconfiguration (SLAAC) [RFC4862] conveys network configuration information in Router Advertisement (RA) messages, to be employed by local hosts for automatic network configuration.

IPv6 largely assumes information stability, with network configuration information changes only taking place in a planned manner: old information is deprecated via reduced lifetimes, while new information is introduced (with longer lifetimes) at the same time. However, there are several scenarios where stale information is not gracefully phased out -- that is, existing information abruptly becomes invalid, while new (replacing) information becomes available. These events, particularly when affecting network prefixes, are commonly referred to as "flash-renumbering events".

In some of these scenarios, the local router producing the network renumbering event may try to deprecate (and eventually invalidate) the currently employed prefix (by explicitly signaling the network about the renumbering event), whereas in other scenarios, it may be unable to do so. In such cases, hosts may end up employing invalid information which results in interoperability problems.

A detailed discussion of this problem and some of the common scenarios where this problem may arise can be found in [RFC8978].

This document updates the Neighbor Discovery specification [RFC4861], the Stateless Address Autoconfiguration (SLAAC) specification [RFC4862], and other associated specifications ([RFC4191], [RFC8106], and [RFC9096]), such that hosts can more gracefully deal with the so-called flash renumbering events [RFC8978], thus improving the robustness of SLAAC.

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

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

Flash renumbering:
A network renumbering event, when an old prefix, used to address hosts, becomes invalid and is replaced by a new prefix. Before the flash renumbering event only the old prefix provides connectivity, and after the flash renumbering only the new one can be used. In other words, there is no period of time when addresses from both prefixes provide connectivity. See [RFC8978] for more detailed discussion of various flash-renumbering scenarios. Note: typically, when flash-renumbering events occur, other IPv6 netowrk configuration information (such as Recursive DNS Server (RDNSS) information [RFC8106]) is affected in the same manner, and thus the term "flash-renumbering" is also employed to refer to a more general "flash-reconfiguration" event.

PIO:
Prefix Information Option, [RFC4861],

RA:
Router Advertisement, [RFC4861].

SLAAC:
IPv6 Stateless Address AutoConfugration, [RFC4862].

SLAAC host:
A host which employs SLAAC for IPv6 network configuration.

SLAAC Router:
A IPv6 router that advertises configuration information via SLAAC.

4. SLAAC reaction to Flash-renumbering Events

In some flash-renumbering scenarios, the local router may try to deprecate the stale information by explicitly signaling the network about the renumbering event, whereas in other scenarios the renumbering event may happen inadvertently, without the router explicitly signaling the scenario to local hosts. The following subsections analyze specific considerations for each of these scenarios.

4.1. Renumbering without Explicit Signaling

In the absence of explicit signalling from SLAAC routers, stale SLAAC configuration information will employed as allowed by the associated lifetimes values. For example, stale prefixes will remain preferred and valid according to the Preferred Lifetime and Valid Lifetime parameters (respectively) of the last received Prefix Information Option (PIO). [RFC4861] specifies the following default values for PIOs:

* Preferred Lifetime (AdvPreferredLifetime): 604800 seconds (7 days)

* Valid Lifetime (AdvValidLifetime): 2592000 seconds (30 days)

This means that, in the absence of explicit signaling by a SLAAC router to deprecate a prefix, it will take a host 7 days (one week) to deprecate the corresponding addresses, and 30 days (one month) to eventually remove any addresses configured for the stale prefix. Clearly, employing such long default values is unacceptable for most deployment scenarios that may experience flash-renumbering events.

NOTE:
[RFC8978] provides an operational recommendation for Customer Edge (CE) routers to override the standard default Preferred Lifetime (AdvPreferredLifetime) and Valid Lifetime (AdvValidLifetime) to 2700 seconds (45 minutes) and 5400 seconds (90 minutes), respectively, thus improving the state of affairs for CE router scenarios.

Similarly, other Neighbor Discovery options employ long default lifetimes that are unacceptable for most deployment scenarios where flash-renumbering events may be experienced.

Use of more appropriate timers in Router Advertisement messages can help limit the amount of time that hosts will maintain stale configuration information. Thus, Section 5.1 specifies more appropriate (i.e., shorter) default lifetimes for Neighbor Discovery options. Section 5.5 provides recommendations about conveying Neighbor Discovery information into RA messages, to help hosts infer when information may have become stale.

4.2. Renumbering with Explicit Signaling

In scenarios where a local router is aware of the renumbering event, it may try to phase out the stale network configuration information. In these scenarios, there are two aspects to be considered:

* The amount of time during which the router should continue trying to deprecate the stale network configuration information.

* The ability of SLAAC hosts to phase out stale configuration.

Since the network could become partitioned at any point in time and for an arbitrarily long period of time, in order to reliably deprecate stale information, a router should try to deprecate such information for its maximum possible lifespan.

NOTE:
For example, a router should try to deprecate a prefix (via a PIO) for a period of time equal to the "Preferred Lifetime" used when advertising the prefix, and try to invalidate the prefix for a period of time equal to the "Valid Lifetime" used when advertising the prefix.

Once the number of seconds in the original "Preferred Lifetime" have elapsed, all hosts will have deprecated the corresponding addresses, while once the number of seconds in the "Valid Lifetime" have elapsed, the corresponding addresses will have been invalidated and removed.

Thus, use of more appropriate default lifetimes for Neighbor Discovery options, as specified in Section 5.1, will reduce the amount of time stale options would need to be advertised by a router to ensure that the associated information is reliably phased out.

In the case of PIOs, in scenarios where a router has positive knowledge that a prefix has become invalid (and thus could signal this condition to local hosts), the current specifications will prevent SLAAC hosts from fully recovering from such stale information: Item "e)" of Section 5.5.3 of [RFC4862] specifies that an RA may never reduce the "RemainingLifetime" to less than 2 hours. Additionally, if the "RemainingLifetime" of an address is less than 2 hours, then a "Valid Lifetime" less than 2 hours will be ignored. The inability to invalidate a stale prefix may prevent communications with the new "owners" of a prefix, and thus is highly undesirable. On the other hand, the Preferred Lifetime of an address may be reduced to any value to avoid the use of addresses from a stale prefix for new communications.

Section 5.4 formally updates [RFC4862] to remove this restriction, such that hosts may react to the advertised "Valid Lifetime" even if it is less than 2 hours. Section 5.3 recommends that routers disseminate network configuration information when a network interface is initialized or reconfigured, such that configuration information propagates in a timelier manner.

5. Improvements to Stateless Address Autoconfiguration (SLAAC)

The following subsections update [RFC4191], [RFC4861], [RFC4862], [RFC8106], and [RFC9096], such that the problem discussed in this document is mitigated. Each of the following subsections improve different aspects of SLAAC, and thus are mostly orthogonal:

* Reduce the default lifetimes of Neighbor Discovery options (Section 5.1):

  This helps limit the amount of time a host may employ stale information, and also limits the amount of time a router should try to deprecate stale information.

* Signal Stale Configuration Information (Section 5.2):

  This allows local hosts to learn about stale configuration information in a timelier manner.

* Honor PIOs with small Valid Lifetimes (Section 5.4):

  This allows hosts to honor PIOs with a Valid Lifetime less than 2 hours, thus resulting in a timelier reaction to flash-renumbering events.

* Recommend routers to retransmit configuration information upon interface initialization/reconfiguration (Section 5.3):

  This helps spread the network configuration information in a timelier manner.

* Recommend routers to always send all options (i.e. the complete configuration information) in RA messages, and in the smallest possible number of packets (Section 5.5):

  This helps propagate the same information to all hosts.

5.1. More Appropriate Neighbor Discovery Option Lifetimes

This document formally updates [RFC4861] to introduce an improved default setting for the MinRtrAdvInterval Neighor Discovery parameter specified in [RFC4861]:

MinRtrAdvInterval= max(3, MaxRtrAdvInterval/3)

RATIONALE:
This expression essentially sets essentially MinRtrAdvInterval to MaxRtrAdvInterval/3, but ensures MinRtrAdvInterval is never smaller than 3 (seconds).

As noted in Appendix A, a number of Neighbor Discovery parameters, such as MinRtrAdvInterval, MaxRtrAdvInterval, preferred lifetimes, and valid lifetimes are related with the link properties and need to have congruent default values and settings.

This document defines the following constants to be employed for the default lifetimes of Neighbor Discovery options:

* ND_DEFAULT_PREFERRED_LIFETIME: 2700 seconds (45 minutes)

* ND_DEFAULT_VALID_LIFETIME: 3600 seconds (60 minutes)

Implementations MAY override these default values according to the considerations in Appendix A.

This document formally updates [RFC4861] to modify the default value of the Router Lifetime field of RA messages as follows:

* AdvDefaultLifetime: ND_DEFAULT_VALID_LIFETIME

NOTE:
This is to align the Router Lifetime with the recommendations in [RFC7772].

This document formally updates [RFC4861] to modify the default values of the Preferred Lifetime (AdvPreferredLifetime) and the Valid Lifetime (AdvValidLifetime) of PIOs as follows:

* AdvPreferredLifetime: ND_DEFAULT_PREFERRED_LIFETIME

* AdvValidLifetime: ND_DEFAULT_VALID_LIFETIME

This document formally updates [RFC4191] to specify the default Route Lifetime of Route Information Options (RIOs) as follows:

* Route Lifetime: It defaults to ND_DEFAULT_VALID_LIFETIME

This document formally updates [RFC8106] to modify the default Lifetime of Recursive DNS Server Options as:

* Lifetime: It defaults to ND_DEFAULT_VALID_LIFETIME

Additionally, this document formally updates [RFC8106] to modify the default Lifetime of DNS Search List Options as:

* Lifetime: It defaults to ND_DEFAULT_VALID_LIFETIME

This document introduces the following update to Section 4 of [RFC9096]:

OLD TEXT:

=======

* ND_PREFERRED_LIMIT: 2700 seconds (45 minutes)

* ND_VALID_LIMIT: 5400 seconds (90 minutes)

RATIONALE:

* These values represent a trade-off among a number of factors, including responsiveness and possible impact on the battery life of connected devices [RFC7772].

* ND_PREFERRED_LIMIT is set according to the recommendations in [RFC7772] for the "Router Lifetime", following the rationale from Section 3.2 of [RFC8978].

* ND_VALID_LIMIT is set to 2 * ND_PREFERRED_LIMIT to provide some additional leeway before configuration information is finally discarded by the hosts.

=======

NEW TEXT:

=======

* ND_PREFERRED_LIMIT: 2700 seconds (45 minutes)

* ND_VALID_LIMIT: 3600 seconds (60 minutes)

=======

NOTE: This aligns the recommended values in [RFC9096] with the default values specified in this section.

5.2. Signaling Stale Configuration Information

In some scenarios, a SLAAC router may learn that previously advertised information has become stale. For example, this may happen when e.g. the advertised information is derived from information that has been dynamically learned from an upstream router via DHCPv6-PD, but the upstream router is no longer in use or available. In such scenarios, it is paramount that the SLAAC router signals the SLAAC configuration information change, to aid hosts in quickly phasing out the stale network configuration information.

SLAAC routers MUST signal stale configuration information by following the guidelines in Section 3.5 ("Signaling Stale Configuration Information") of [RFC9096].

In scenarios where flash renumbering events or configuration changes are frequent, a router may end up in a situation where multiple pieces of information may need to be simultaneously deprecated, and thus the size of Router Advertisement messages could substantially increase. In such scenarios, routers MAY limit themselves to deprecate the most recent configuration that would fit into a single Router Advertisement message without fragmentation.

5.3. Propagating Interface Configuration Changes

When the information to be contained in RAs changes (e.g. an interface is reconfigured), it is paramount that updated information is propagated to hosts connected to the corresponding network in a timely manner. Thus, this document replaces the following text from Section 6.2.4 of [RFC4861]:

> The information contained in Router Advertisements may change through actions of system management. For instance, the lifetime of advertised prefixes may change, new prefixes could be added, a router could cease to be a router (i.e., switch from being a router to being a host), etc. In such cases, the router MAY transmit up to MAX_INITIAL_RTR_ADVERTISEMENTS unsolicited advertisements, using the same rules as when an interface becomes an advertising interface.

with:

> The information contained in Router Advertisements may change through actions of system management, or because it is derived from information learned from an upstream router (via e.g. DHCPv6 [RFC8415]), and such information has changed. For instance, the lifetime of advertised prefixes may change, new prefixes could be added, existing prefixes could be removed, a router could cease to be a router (i.e., switch from being a router to being a host), etc. In such cases, the router MUST transmit MAX_INITIAL_RTR_ADVERTISEMENTS unsolicited advertisements, using the same rules as when an interface becomes an advertising interface.

RATIONALE:

* Use of stale information can lead to interoperability problems. Therefore, it is important that new configuration information propagates in a timelier manner to all hosts.

5.4. Honor Small PIO Valid Lifetimes

This document introduces the following update to Section 5.5.3 of [RFC4862]:

OLD TEXT:

=======

e) If the advertised prefix is equal to the prefix of an address configured by stateless autoconfiguration in the list, the preferred lifetime of the address is reset to the Preferred Lifetime in the received advertisement. The specific action to perform for the valid lifetime of the address depends on the Valid Lifetime in the received advertisement and the remaining time to the valid lifetime expiration of the previously autoconfigured address. We call the remaining time "RemainingLifetime" in the following discussion:

1. If the received Valid Lifetime is greater than 2 hours or greater than RemainingLifetime, set the valid lifetime of the corresponding address to the advertised Valid Lifetime.

2. If RemainingLifetime is less than or equal to 2 hours, ignore the Prefix Information option with regards to the valid lifetime, unless the Router Advertisement from which this option was obtained has been authenticated (e.g., via Secure Neighbor Discovery [RFC3971]). If the Router Advertisement was authenticated, the valid lifetime of the corresponding address should be set to the Valid Lifetime in the received option.

3. Otherwise, reset the valid lifetime of the corresponding address to 2 hours.

The above rules address a specific denial-of-service attack in which a bogus advertisement could contain prefixes with very small Valid Lifetimes. Without the above rules, a single unauthenticated advertisement containing bogus Prefix Information options with short Valid Lifetimes could cause all of a node’s addresses to expire prematurely. The above rules ensure that legitimate advertisements (which are sent periodically) will "cancel" the short Valid Lifetimes before they actually take effect.

Note that the preferred lifetime of the corresponding address is always reset to the Preferred Lifetime in the received Prefix Information option, regardless of whether the valid lifetime is also reset or ignored. The difference comes from the fact that the possible attack for the preferred lifetime is relatively minor. Additionally, it is even undesirable to ignore the preferred lifetime when a valid administrator wants to deprecate a particular address by sending a short preferred lifetime (and the valid lifetime is ignored by accident).

=======

NEW TEXT:

=======

e) If the advertised prefix is equal to the prefix of an address configured by stateless autoconfiguration in the list, the valid lifetime and the preferred lifetime of the address should be updated by processing the Valid Lifetime and the Preferred Lifetime (respectively) in the received advertisement.

While allowing updates to the valid lifetime in RAs could enable an attacker to invalidate addresses by setting the valid lifetime to zero, this does not significantly worsen the security situation. An attacker capable of sending rogue RAs already has the power to disrupt connectivity by manipulating other parameters, such as gateway or DNS information. Therefore, accepting a zero lifetime does not make the system more vulnerable than it already is, as invalidating the prefix is just one of the many vectors available to perform DoS attacks to on-link node (see [RFC3756]).

In scenarios where RA-based attacks are of concern, mitigations such as RA-Guard [RFC6105] [RFC7113] or SEND [RFC3971] should be implemented.

=======

RATIONALE:

* This change allows hosts to react to the signal provided by a router that has positive knowledge that a prefix is not assigned to the given link anymore. In particular it would allow the host to invalidate addresses from that prefix, and, if the prefix is reassigned to another link, allows the host to communicate to devices on that link.

* The behavior described in [RFC4862] had been incorporated during the revision of the original IPv6 Stateless Address Autoconfiguration specification ([RFC1971]). At the time, the IPNG working group decided to mitigate the attack vector represented by Prefix Information Options with very short lifetimes, on the premise that these packets represented a bigger risk than other ND-based attack vectors [IPNG-minutes].

  While reconsidering the trade-offs represented by such decision, we conclude that the drawbacks of the aforementioned mitigation outweigh the possible benefits, as specified in the updated text.

5.5. Conveying Information in Router Advertisement (RA) Messages

Intentionally omitting information in Router Advertisements may prevent the propagation of such information, and may represent a challenge for hosts that need to infer whether they have received a complete set of SLAAC configuration information. As a result, this section recommends that, to the extent that is possible, RA messages contain a complete set of SLAAC information.

This document replaces the following text from Section 6.2.3 of [RFC4861]:

> A router MAY choose not to include some or all options when sending unsolicited Router Advertisements. For example, if prefix lifetimes are much longer than AdvDefaultLifetime, including them every few advertisements may be sufficient. However, when responding to a Router Solicitation or while sending the first few initial unsolicited advertisements, a router SHOULD include all options so that all information (e.g., prefixes) is propagated quickly during system initialization.
>
> If including all options causes the size of an advertisement to exceed the link MTU, multiple advertisements can be sent, each containing a subset of the options.

with:

> A router SHOULD include all options in a single Router Advertisement. However, there are scenarios when routers MAY split the information between multiple RAs. In particular:
>
> * Routers MAY be explicitly or implicitly configured to send multiple RAs and split information between them. For example, a router could be configured to send information associated with different provisioning domains [RFC7556] in different RAs, or to send multiple RAs, one per VRRPv3 [RFC9568] group.
>
> * If including all options causes the size of an RA to exceed the link MTU, multiple RAs SHOULD be sent, each containing a subset of the options. Routers SHOULD whenever possible, split the information between the fewest possible number of RAs.

RATIONALE:

* Sending information in the smallest possible number of packets was somewhat already implied by the original text in [RFC4861]. Including all options when sending RAs leads to simpler code (as opposed to dealing with special cases where specific information is intentionally omitted), helps hosts infer when

they have received a complete set of SLAAC configuration information, and reduces the probability of hosts learning only a partial subset of SLAAC configuration information. Note that while [RFC4861] allowed some RAs to omit some options, to the best of the authors' knowledge, all SLAAC router implementations always send all options in the smallest possible number of packets. Therefore, this section simply aligns the protocol specifications with existing implementation practice.

* However in some scenarios (including, but not limited to multihoming or having a router providing information from multiple configuration or provisional domains (PvD) to non-PvD-aware hosts) it might be desirable to send multiple sets of network configuration information in multiple RAs.

6. IANA Considerations

This document has no actions for IANA.

7. Implementation Status

[NOTE: This section is to be removed by the RFC-Editor before this document is published as an RFC.]

This section summarizes the implementation status of the updates proposed in this document. In some cases, they correspond to variants of the mitigations proposed in this document (e.g., use of reduced default lifetimes for PIOs, albeit using different values than those recommended in this document). In such cases, we believe these implementations signal the intent to deal with the problems described in [RFC8978] while lacking any guidance on the best possible approach to do it.

7.1. More Appropriate Lifetime Values

7.1.1. Router Configuration Variables

7.1.1.1. rad(8)

We have produced a patch for OpenBSD's rad(8) [rad] that employs reduced lifetimes for Neighbor Discovery options, as recommended in this document. The patch is available at: <https://www.gont.com.ar/code/fgont-patch-rad-pio-lifetimes.txt>.

7.1.1.2. radvd(8)

The radvd(8) daemon [radvd], normally employed by Linux-based router implementations, currently employs different default lifetimes than those recommended in [RFC4861]. radvd(8) employs the following default values [radvd.conf]:

* Preferred Lifetime: 14400 seconds (4 hours)

* Valid Lifetime: 86400 seconds (1 day)

These values do not follow the recommendations in this document, but nevertheless represent a deviation and improvement from the current standards.

7.2. Honor Small PIO Valid Lifetimes

7.2.1. Linux Kernel

A Linux kernel implementation ohas been committed to the net-next tree. The implementation was produced in April 2020 by Fernando Gont <fgont@si6networks.com>. The corresponding patch can be found at: <https://patchwork.ozlabs.org/project/netdev/patch/20200419122457.GA971@archlinux-current.localdomain/>

7.2.2. NetworkManager

NetworkManager [NetworkManager] processes RA messages with a Valid Lifetime less than 2 hours as recommended in this document.

7.3. Conveying Information in Router Advertisement (RA) Messages

We know of no implementation that splits network configuration information into multiple RA messages.

7.4. Recovery from Stale Configuration Information without Explicit Signaling

7.4.1. dhcpcd(8)

The dhcpcd(8) daemon [dhcpcd], a user-space SLAAC implementation employed by some Linux-based and BSD-derived operating systems, will set the Preferred Lifetime of addresses corresponding to a given prefix to 0 when a single RA from the router that previously advertised the prefix fails to advertise the corresponding prefix. However, it does not affect the corresponding Valid Lifetime. Therefore, it can be considered a partial implementation of this feature.

7.5. Other mitigations implemented in products

[FRITZ] is a Customer Edge Router that tries to deprecate stale prefixes by advertising stale prefixes with a Preferred Lifetime of 0, and a Valid Lifetime of 2 hours (or less). There are two things to note with respect to this implementation:

* Rather than recording prefixes on stable storage (as recommended in [RFC9096]), this implementation checks the source address of IPv6 packets, and assumes that usage of any address that does not correspond to a prefix currently-advertised by the Customer Edge Router is the result of stale network configuration information. Hence, upon receipt of a packet that employs a source address that does not correspond to a currently-advertised prefix, this implementation will start advertising the corresponding prefix with small lifetimes, with the intent of deprecating it.

* Possibly as a result of item "e)" (pp. 19-20) from Section 5.5.3 of [RFC4862] (discussed in Section 5.4 of this document), upon first occurrence of a stale prefix, this implementation will employ a decreasing Valid Lifetime, starting from 2 hours (7200 seconds), as opposed to a Valid Lifetime of 0.

8. Security Considerations

The protocol update in Section 5.4 could allow an on-link attacker to perform a Denial of Service attack against local hosts, by sending a forged RA with a PIO with a Valid Lifetime of 0. Upon receipt of that packet, local hosts would invalidate the corresponding prefix, and therefore remove any addresses configured for that prefix, possibly terminating e.g. associated TCP connections. However, an attacker may achieve similar effects via a number other Neighbor Discovery (ND) attack vectors, such as directing traffic to a non-existing node until ongoing TCP connections time out, or performing a ND-based man-in-the-middle (MITM) attack and subsequently forging TCP RST segments to cause on-going TCP connections to be reset. Thus, for all practical purposes, this attack vector does not really represent any greater risk than other ND attack vectors. As noted in Section 5.4 , in scenarios where RA-based attacks are of concern, proper mitigations such as RA-Guard [RFC6105] [RFC7113] or SEND [RFC3971] should be implemented.

9. Acknowledgments

The authors would like to thank (in alphabetical order) Mikael Abrahamsson, Tore Anderson, Luis Balbinot, Brian Carpenter, Lorenzo Colitti, Owen DeLong, Gert Doering, Thomas Haller, Nick Hilliard, Bob Hinden, Philip Homburg, Lee Howard, Christian Huitema, Tatuya Jinmei, Erik Kline, Ted Lemon,Albert Manfredi, Roy Marples, Florian Obser, Jordi Palet Martinez, Michael Richardson, Hiroki Sato, Mark Smith, Hannes Frederic Sowa, Dave Thaler, Tarko Tikan, Ole Troan, Eduard Vasilenko, and Loganaden Velvindron, for providing valuable comments on earlier versions of this document.

Fernando would like to thank Alejandro D'Egidio and Sander Steffann for a discussion of these issues, which led to the publication of [RFC8978], and eventually to this document.

Fernando would also like to thank Brian Carpenter who, over the years, has answered many questions and provided valuable comments that has benefited his protocol-related work.

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>.

[RFC7556] Anipko, D., Ed., "Multiple Provisioning Domain Architecture", RFC 7556, DOI 10.17487/RFC7556, June 2015, <https://www.rfc-editor.org/info/rfc7556>.

[RFC7772] Yourtchenko, A. and L. Colitti, "Reducing Energy Consumption of Router Advertisements", BCP 202, RFC 7772, DOI 10.17487/RFC7772, February 2016, <https://www.rfc-editor.org/info/rfc7772>.

[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>.

10.2. Informative References

[dhcpcd] Marples, R., "dhcpcd - a DHCP client", <https://roy.marples.name/projects/dhcpcd/>.

[FRITZ] Gont, F., "Quiz: Weird IPv6 Traffic on the Local Network (updated with solution)", SI6 Networks Blog, February 2016, <https://www.si6networks.com/2016/02/16/quiz-weird-ipv6-traffic-on-the-local-network-updated-with-solution/>.

[IPNG-minutes] IETF, "IPNG working group (ipngwg) Meeting Minutes", Proceedings of the thirty-eightt Internet Engineering Task Force , April 1997, <https://www.ietf.org/proceedings/38/97apr-final/xrtftr47.htm>.

[NetworkManager] NetworkManager, "NetworkManager web site", <https://wiki.gnome.org/Projects/NetworkManager>.

[rad] Obser, F., "OpenBSD Router Advertisement Daemon - rad(8)", <https://cvsweb.openbsd.org/src/usr.sbin/rad/>.

[radvd] Hawkins, R. and R. Johnson, "Linux IPv6 Router Advertisement Daemon (radvd)", <http://www.litech.org/radvd/>.

[radvd.conf] Hawkins, R. and R. Johnson, "radvd.conf - configuration file of the router advertisement daemon", <https://github.com/reubenhwk/radvd/blob/master/radvd.conf.5.man>.

[RFC1971] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 1971, DOI 10.17487/RFC1971, August 1996, <https://www.rfc-editor.org/info/rfc1971>.

[RFC3756] Nikander, P., Ed., Kempf, J., and E. Nordmark, "IPv6 Neighbor Discovery (ND) Trust Models and Threats", RFC 3756, DOI 10.17487/RFC3756, May 2004, <https://www.rfc-editor.org/info/rfc3756>.

[RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, DOI 10.17487/RFC3971, March 2005, <https://www.rfc-editor.org/info/rfc3971>.

[RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", RFC 4191, DOI 10.17487/RFC4191, November 2005, <https://www.rfc-editor.org/info/rfc4191>.

[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>.

[RFC7113] Gont, F., "Implementation Advice for IPv6 Router Advertisement Guard (RA-Guard)", RFC 7113, DOI 10.17487/RFC7113, February 2014, <https://www.rfc-editor.org/info/rfc7113>.

[RFC8106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", RFC 8106, DOI 10.17487/RFC8106, March 2017, <https://www.rfc-editor.org/info/rfc8106>.

[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>.

[RFC8978] Gont, F., 貼or転, J., and R. Patterson, "Reaction of IPv6 Stateless Address Autoconfiguration (SLAAC) to Flash-Renumbering Events", RFC 8978, DOI 10.17487/RFC8978, March 2021, <https://www.rfc-editor.org/info/rfc8978>.

[RFC9096] Gont, F., 貼or転, J., Patterson, R., and B. Volz, "Improving the Reaction of Customer Edge Routers to IPv6 Renumbering Events", BCP 234, RFC 9096, DOI 10.17487/RFC9096, August 2021, <https://www.rfc-editor.org/info/rfc9096>.

[RFC9568] Lindem, A. and A. Dogra, "Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6", RFC 9568, DOI 10.17487/RFC9568, April 2024, <https://www.rfc-editor.org/info/rfc9568>.

Appendix A. Selecting Neighbor Discovery Lifetimes

While many default values in from the Neighbor Discovery specification [RFC4861] assume fairly reliable communication of Neighbor Discovery messages, communication of multicasted RA messages tends to be rather unreliable for battery-powered devices, which tend to drop many of such messages to reduce the associated effects on power consumption [RFC7772]. The expressions in this section may be employed to override the default lifetime values from Section 5.1 while considering packet loss.

The following relationship exists among Neighbor Discovery parameters:

ND_PREFERRED_LIFETIME= ND_RAS_PREFERRED * MaxRtrAdvInterval

ND_VALID_LIFETIME= ND_RAS_VALID * MaxRtrAdvInterval

where:

ND_PREFERRED_LIFETIME:
: Preferred lifetime for Neighbor Discovery information (where applicable). This parameter is the value that would be employed to override the ND_DEFAULT_PREFERRED_LIFETIME value specified in Section 5.1 of this document.

ND_VALID_LIFETIME:
: Valid lifetime for Neighbor Discovery information (where applicable). This parameter is the value that would be employed to override the ND_DEFAULT_VALID_LIFETIME value specified in Section 5.1 of this document.

ND_RAS_PREFERRED:
: Number of RA messages sent during the ND_PREFERRED_LIFETIME period..

ND_RAS_VALID:
: Number of RA messages sent during the ND_VALID_LIFETIME.

MaxRtrAdvInterval:
: Maximum time allowed between sending unsolicited multicast Router Advertisements from the interface, in seconds (as specified in [RFC4861]. It defaults to 300 seconds.

ND_RAS_PREFERRED and ND_RAS_VALID should be computed with the expression:

$$n >= \ln(1 - P)/\ln(Loss)$$

where "n" is the number of RA messages that a router should send, such that, given an RA-message loss rate of "Loss", there is a probability of "P" that at least one of such messages is received by the target hosts.

NOTES:

As noted in Section 6.2.4 of [RFC4861], RA messages are retransmitted with uniformly distributed random interval between the interface’s configured MinRtrAdvInterval and MaxRtrAdvInterval. Thus, the equation above represents the worst-case scenario, where each RA message is retransmitted at MaxRtrAdvInterval seconds.

It should be evident from the previous expressions that for any given packet loss ("Loss") and probability "P", ND_RAS_PREFERRED and ND_RAS_VALID express the relationship of the Preferred Lifetime (ND_PREFERRED_LIFETIME) and the Valid Lifetime (ND_VALID_LIFETIME) with the sending rate (as derived from MaxRtrAdvInterval). Therefore, if e.g. the ND_PREFERRED_LIFETIME or ND_VALID_LIFETIME are reduced, MaxRtrAdvInterval should be reduced accordingly such that the probability "P" is not affected.

The following table tabulates the value of P (probability of receiving at least one RA message) for a combination of "n" (number of RA messages sent) and Loss (Loss rate for multicasted RA messages):

n / Loss	0.10	0.20	0.30	0.40	0.50	0.60	0.70	0.80	0.90	0.95
3	0.99900	0.99200	0.97300	0.93600	0.87500	0.78400	0.65700	0.48800	0.27100	0.14263
4	0.99990	0.99840	0.99190	0.97440	0.93750	0.87040	0.75990	0.59040	0.34390	0.18549
5	0.99999	0.99968	0.99757	0.98976	0.96875	0.92224	0.83193	0.67232	0.40951	0.46856
6	1.00000	0.99994	0.99927	0.99590	0.98437	0.95334	0.88235	0.73786	0.46856	0.26491
7	1.00000	0.99999	0.99978	0.99836	0.99219	0.97201	0.91765	0.79028	0.52170	0.30166
8	1.00000	1.00000	0.99993	0.99934	0.99609	0.98320	0.94235	0.83223	0.56953	0.33658
9	1.00000	1.00000	0.99998	0.99974	0.99805	0.98992	0.95965	0.86578	0.61258	0.36975
10	1.00000	1.00000	0.99999	0.99989	0.99902	0.99395	0.97175	0.89263	0.65132	0.40126
11	1.00000	1.00000	0.99999	0.99995	0.99951	0.99637	0.98022	0.91410	0.68618	0.43119
12	1.00000	1.00000	0.99999	0.99998	0.99975	0.99782	0.98615	0.93128	0.71757	0.45963

Table 1: Sample values for P = 1 - (Loss)^n

Appendix B. Rationale for the default values specified in this document

The default values from Section 5.1 result, when employing the expressions from Appendix A, in the following values:

ND_RAS_PREFERRED= 9

ND_RAS_VALID= 12

We note that for e.g. for an RA loss rate of 50% (Loss=0.50), this would result in a probability of hosts refreshing this timer before it expires of 0.99805. We note that if the Preferred Lifetime expires, and the host has configured addresses for other prefixes, it will start preferring those other addresses instead. On the other hand, if the host has not configured addresses for other prefixes, it may still employ addresses even if they are not "Preferred" (please see Section 5.5.4 of [RFC4862]). Similarly, for the same loss rate of 50% (Loss=0.50), this would result in a probability of hosts refreshing this timer before it expires of 0.99975.

Authors' Addresses

Fernando Gont
SI6 Networks
Segurola y Habana 4310, 7mo Piso
Villa Devoto
Ciudad Autonoma de Buenos Aires
Argentina
Email: fgont@si6networks.com
URI: https://www.si6networks.com

Jan Zorz
6connect
Email: jan@6connect.com

Richard Patterson
Sky UK
Email: richard.patterson@sky.uk

Jen Linkova (editor)
Google
Email: furry13@gmail.com, furry@google.com