

IPv6 Maintenance (6man) Working Group
Internet-Draft
Updates: [4861](#), [4862](#) (if approved)
Intended status: Standards Track
Expires: September 10, 2020

F. Gont
SI6 Networks / UTN-FRH
J. Zorz
Go6 Institute
R. Patterson
Sky UK
March 9, 2020

Improving the Robustness of Stateless Address Autoconfiguration (SLAAC)
to Flash Renumbering Events
[draft-gont-6man-slaac-renum-03](#)

Abstract

In renumbering scenarios where an IPv6 prefix suddenly becomes invalid, hosts on the local network will continue using stale prefixes for an unacceptably long period of time, thus resulting in connectivity problems. This document improves the reaction of IPv6 Stateless Address Autoconfiguration to such renumbering scenarios.

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 September 10, 2020.

Copyright Notice

Copyright (c) 2020 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 | 2 |
| 2. | Terminology | 3 |
| 3. | SLAAC reaction to Flash-renumbering Events | 3 |
| 3.1. | Renumbering without Explicit Signaling | 3 |
| 3.2. | Renumbering with Explicit Signaling | 4 |
| 4. | Improvements to Stateless Address Autoconfiguration (SLAAC) . | 5 |
| 4.1. | More Appropriate Lifetime Values | 6 |
| 4.1.1. | Router Configuration Variables | 6 |
| 4.1.2. | Processing of PIO Lifetimes at Hosts | 7 |
| 4.2. | Processing of PIO lifetimes | 8 |
| 4.3. | Interface Initialization | 8 |
| 4.4. | Conveying Information in Router Advertisement (RA) Messages | 9 |
| 4.5. | Recovery from Stale Configuration Information without Explicit Signaling | 10 |
| 5. | IANA Considerations | 13 |
| 6. | Security Considerations | 13 |
| 7. | Acknowledgments | 14 |
| 8. | References | 14 |
| 8.1. | Normative References | 14 |
| 8.2. | Informative References | 15 |
| Appendix A. | Sample Timeline for Host Processing of RAs | 16 |
| Appendix B. | Analysis of Some Suggested Workarounds | 17 |
| B.1. | On a Possible Reaction to ICMPv6 Error Messages | 18 |
| B.2. | On a Possible Improvement to Source Address Selection . . | 18 |
| | Authors' Addresses | 20 |

[1.](#) Introduction

IPv6 network renumbering is expected to take place in a planned manner, with old/stale prefixes being phased-out via reduced prefix lifetimes while new prefixes (with normal lifetimes) are introduced. However, there are a number of scenarios that may lead to the so-called "flash-renumbering" events, where the prefix being employed on a network suddenly becomes invalid and replaced by a new prefix [[I-D.ietf-v6ops-slaac-renum](#)]. In such scenarios, hosts on the local network will continue using stale prefixes for an unacceptably long period of time, thus resulting in connectivity problems. [[I-D.ietf-v6ops-slaac-renum](#)] discusses this problem in detail.

In some scenarios, the local router producing the network renumbering event may try to deprecate the currently-employed prefixes (thus explicitly signaling the network about the renumbering event), whereas in other scenarios it may be unaware about the renumbering event and thus unable signal hosts about it.

From the perspective of a Stateless Address Autoconfiguration (SLAAC) host, there are two different (but related) problems to be solved:

- o Avoiding the use of stale addresses for new communication instances
- o Performing "garbage collection" for the stale prefixes (and related network configuration information)

Clearly, if a host has both working and stale addresses, it is paramount that it employs working addresses for new communication instances. Additionally, a host should also perform garbage collection for the stale prefixes/addresses, since they not only tie system resources, but also prevent communication with the new "owners" of the stale prefixes.

2. Terminology

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. SLAAC reaction to Flash-renumbering Events

As noted in [Section 1](#), in some scenarios the router triggering the renumbering event may be able to explicitly signal the network about this event, while in other scenarios the renumbered hosts may need to infer a renumbering event is taking place. The following subsections analyze specific considerations for each of these scenarios.

3.1. Renumbering without Explicit Signaling

In the absence of explicit signalling from SLAAC routers (such as sending Prefix Information Options (PIOs) with small lifetimes to deprecate the stale prefixes), stale prefixes will remain preferred and valid according to the Preferred Lifetime and Valid Lifetime values (respectively) of the last received PIO. IPv6 SLAAC employs the following default values for PIOs:

- o Preferred Lifetime (AdvPreferredLifetime): 604800 seconds (7 days)
- o 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 un-prefer the corresponding addresses, and 30 days (one month) to eventually remove any addresses configured for the stale prefix. Clearly, for any practical purposes, employing such long default values are the equivalent of not using any timers at all, since taking 7 days or 30 days (respectively) to recover from a network problem is simply unacceptable.

Use of more appropriate timers in Router Advertisement messages can help limit the amount of time that hosts will maintain stale configuration information. Additionally, hosts are normally in a position to infer that a prefix has become invalid -- for example, if a given router ceases to advertise an existing prefix and at the same time starts to advertise a new prefix.

[Section 4.1.1](#) recommends the use of more appropriate lifetimes for PIOs, while [Section 4.1.2](#) proposes to cap the accepted Valid Lifetime and Preferred Lifetime values at hosts, such that more appropriate values are employed even in the presence of legacy routers.

[Section 4.5](#) specifies a local policy that SLAAC hosts can implement to heuristically infer that network configuration information has changed, such that stale configuration can be phased out.

[3.2.](#) Renumbering with Explicit Signaling

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

- o The amount of time during which the router should continue trying to deprecate the stale network configuration information
- o The ability of SLAAC hosts to phase out stale configuration in a timelier manner.

In the absence of Router Advertisements (RAs) that include PIOs that would reduce the Valid Lifetime and Preferred Lifetime of a prefix, hosts would normally employ the lifetime values from PIO option of the last received RA messages. Since the network could be partitioned for an arbitrarily long period of time, a router would need to try to "unprefer" a prefix for the amount of time employed

for the "Preferred Lifetime", and try to invalidate the prefix for the amount of time employed for the "Valid Lifetime" (see [Section 12 of \[RFC4861\]](#)).

NOTE:

Once the number of seconds in the original "Preferred Lifetime" have elapsed, all hosts would have "unpreferred" the corresponding addresses anyway, while once the number of seconds in the "Valid Lifetime" have elapsed, the corresponding addresses would be invalidated and removed.

Thus, use of more appropriate default lifetimes for PIOs, as proposed in [Section 4.1.1](#), would reduce the amount of time a stale prefix would need to be announced as such by a router to unprefer/invalidate it.

In scenarios where a router has positive knowledge that a prefix has become invalid and thus 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" more than to two hours. If the RemainingLifetime of an address is smaller than 2 hours, then a Valid Lifetime smaller than 2 hours will be ignored. The inability to invalidate a stale prefix would prevent communication with the new "owners" of the stale prefix, and thus is highly undesirable. On the other hand, the Preferred Lifetime of an address *can* be reduced to any value to avoid the use of a stale prefix to be employed for new communications.

[Section 4.2](#) updates [\[RFC4862\]](#) such that this restriction is removed, and hosts react to the advertised "Valid Lifetime" (even if it is smaller than 2 hours).

Finally, [Section 4.3](#) recommends that routers disseminate network configuration information when a network interface is initialized, such that possibly new configuration information propagates in a timelier manner.

4. Improvements to Stateless Address Autoconfiguration (SLAAC)

The following subsections update [\[RFC4861\]](#) and [\[RFC4862\]](#), such that the problem discussed in this document is mitigated. The aforementioned updates are mostly orthogonal, and mitigate different aspects of SLAAC that prevent a timely reaction to flash renumbering events.

- o Reduce default Valid Lifetime and Preferred Lifetime of PIOs ([Section 4.1.1](#)):

This helps limit the amount of time a host will employ stale information, and also limits the amount of time a router needs to insist in obsoleting stale information.

- o Cap received Valid Lifetime and Preferred Lifetime of PIOs ([Section 4.1.2](#)):
This helps limit the amount of time a host will employ stale information, even in the presence of legacy ([RFC4861](#)) routers.
- o Honor PIOs with small Valid Lifetimes ([Section 4.2](#)):
This allows routers to invalidate stale prefixes ([RFC4861](#)) prevented hosts from honoring PIOs with a Valid Lifetime smaller than two hours).
- o Recommend routers to retransmit configuration information upon interface initialization/reinitialization ([Section 4.3](#)):
This helps spread the new information and also deprecate stale information via host-side heuristics (see [Section 4.5](#)).
- o 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 4.4](#)):
This helps spread information, and also allows hosts to infer that information missing in RA messages has become stale (see [Section 4.5](#)).
- o Infer stale network configuration information from received RAs ([Section 4.5](#)):
This allows hosts to get rid of stale network configuration information, even in the absence of explicit signaling.

[4.1.](#) More Appropriate Lifetime Values

[4.1.1.](#) Router Configuration Variables

The default value for the "lifetime" parameters in PIOs is updated as follows:

AdvValidLifetime: 48 * AdvDefaultLifetime

AdvPreferredLifetime: AdvDefaultLifetime

NOTE:

[RFC4861](#) specifies AdvDefaultLifetime as 3 * MaxRtrAdvInterval (which defaults to 600 seconds). This means this document specifies AdvPreferredLifetime as 1800 seconds. This document specifies AdvValidLifetime as 48 * AdvDefaultLifetime, resulting in a AdvValidLifetime of 86400 seconds (1 day).

RATIONALE:

- * The default values for PIO lifetimes should be such that, under normal circumstances (including some packet loss), the associated timers are refreshed/reset, but in the presence of network failures (such as network configuration information becoming stale), some fault recovering action (such as un-prefering the corresponding addresses and subsequently removing them) is triggered.
- * In the context of [[RFC8028](#)], where it is clear that the use of addresses configured for a given prefix is tied to the next-hop router that advertised the prefix, the "Preferred Lifetime" of a PIO should never be larger than the "Router Lifetime" (AdvDefaultLifetime) of Router Advertisement messages. Some leeway should be provided for the "Valid Lifetime" to cope with transient network problems.
- * As a result, this document updates [[RFC4861](#)] such that the default Valid Lifetime (AdvValidLifetime) and the default Preferred Lifetime (AdvPreferredLifetime) of PIOs are specified as a function of the default "Router Lifetime" (AdvDefaultLifetime) of Router Advertisement messages.
- * In the absence of RAs that refresh information, addresses configured for advertised prefixes become un-preferred in a timelier manner, and thus Rule 3 of [[RFC6724](#)] will cause other configured addresses (if available) to be preferred.

4.1.2. Processing of PIO Lifetimes at Hosts

Hosts SHOULD cap the "Valid Lifetime" and "Preferred Lifetime" of PIOs as follows:

- o Valid Lifetime= MIN(Valid Lifetime, 48 * "Router Lifetime")
- o Preferred Lifetime= MIN(Preferred Lifetime, "Router Lifetime")

RATIONALE:

- * Capping the lifetimes in PIOs as suggested will not eliminate the problem discussed in this document, but will certainly reduce the amount of time it takes for hosts to converge to updated network configuration information, even when the SLAAC router advertises PIOs with the default values specified in [[RFC4861](#)] (as opposed to the new default values specified in [Section 4.1.1](#)) and when the corresponding router ceases to send RAs.

4.2. Processing of PIO lifetimes

The entire item "e)" (pp. 19-20) from [Section 5.5.3 of \[RFC4862\]](#) is replaced with the following 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 is reset to the Valid Lifetime and the Preferred Lifetime (respectively) in the received advertisement.

RATIONALE:

- * This change allows hosts to react to the information provided by a router that has positive knowledge that a prefix has become invalid.
- * Attacks aiming at disabling an advertised prefix via a Valid Lifetime of 0 are not really more harmful than other attacks that can be performed via forged RA messages, such as those aiming at completely disabling a next-hop router via an RA that advertises a Router Lifetime of 0, or performing a Denial of Service (DoS) attack by advertising illegitimate prefixes via PIOs. In scenarios where RA-based attacks are of concern, proper mitigations such as RA-Guard [[RFC6105](#)] [[RFC7113](#)] should be implemented.

4.3. Interface Initialization

When an interface is initialized, it is paramount that network configuration information is spread on the corresponding network (particularly in scenarios where an interface has been re-initialized, and the conveyed information has changed). Thus, this document replaces the following text from [Section 6.2.4 of \[RFC4861\]](#):

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:

In such cases, the router SHOULD 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 paramount that new configuration information is delivered in a timely manner to all hosts.

NOTE:

[[I-D.ietf-v6ops-cpe-slaac-renum](#)] specifies recommendations for CPE routers, including which information should be included in RA messages to deprecate stale network configuration information (if any).

4.4. Conveying Information in Router Advertisement (RA) Messages

Intentionally omitting information in Router Advertisements helps prevent the propagation of such information. To the best of the authors' knowledge, SLAAC routers always send all options in the smallest possible number of packets, so this section simply more clearly aligns the existing specifications with real implementations.

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:

When sending Router Advertisements, a router SHOULD include all options.

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. In all cases, routers SHOULD convey all information using the smallest possible number of packets.

RATIONALE:

- * Sending information in the smallest possible number of packets was somewhat already implied from the original text in [\[RFC4861\]](#), and in this respect the proposed update just adds clarity. Including all options when sending RAs both leads to simpler code (as opposed to dealing with special cases where specific information is intentionally omitted), and also helps hosts infer of network configuration changes in a timelier manner. Note that while [\[RFC4861\]](#) allowed some RAs to omit some options, the authors of this document know of no implementation of such behavior. Therefore, the proposed change simply reflects normal practice.

4.5. Recovery from Stale Configuration Information without Explicit Signaling

The goal of the algorithm specified in this section is to allow hosts to infer when a previously-advertised prefix has become stale, such that previously-configured addresses are "phased-out" and the host can transition to the newly-advertised prefixes in a timelier manner. Most of the value of this algorithm is in being able to mitigate the problem discussed in this document at hosts themselves, without requiring updates on local routers.

Host can normally infer when network configuration information has changed. For example, if a SLAAC router (as identified by its link-local address) has ceased to advertise a previously-advertised prefix and has also started to advertise new prefixes via PIOs, this should be considered an indication that network configuration information has changed. Implementation of this kind of heuristics would allow a timelier reaction to network configuration changes even in scenarios where there is no explicit signaling from the network -- thus improving robustness.

The basic premise behind this algorithm is that, when a router advertises new prefixes for address configuration (PIO with the "A" bit set), but fails to advertise the previously-advertised prefixes, this is an indication that the previously-advertised prefixes have become stale. Therefore, if this was the only router advertising this prefix, configured addresses for the stale prefixes should be "un-preferred" (such that they are not employed for new communication instances), and they should eventually be removed (if this condition persists).

The algorithm specified in this section updates the state of a configured address upon receipt of a number of consecutive RAs that, while carrying PIOs, fail to advertise a previously-advertised prefix. This algorithm can accommodate the (theoretical) scenario where a router may split PIOs among a number of RA messages.

Local information maintained for each prefix advertised by each router is augmented with one counter named "LTA" (Lifetime Avoidance) that defaults to 0, that counts the number of consecutive RAs received from the corresponding router that do not advertise the corresponding prefix.

NOTE:

Hosts are already expected to keep track of which router has advertised which prefix in order to be able to properly select the first-hop router in multiple-prefix networks [[RFC8028](#)] [[RFC8504](#)]. Throughout this specification, each router is identified by its link-local address.

This algorithm employs two configuration variables:

LTA_RAS_UNPREFER: Number of consecutive RAs *not* carrying a given prefix from a given router that will cause the prefix to become unpreferred/deprecated. It defaults to LTA_RAS_UNPREFER_DEFAULT, which this document specifies as 2.

LTA_RAS_INVALID: Number of consecutive RAs *not* carrying a given prefix from a given router that will cause the prefix to become invalid. It defaults to LTA_RAS_INVALID_DEFAULT, which this document specifies as 4.

After normal processing of Router Advertisement messages, Router Advertisements that contain at least one PIO MUST be processed as follows:

- o The LTA counter for each prefix advertised in the current Router Advertisement, and associated with this particular router, should be set to 0.
- o For each prefix that had been previously advertised by this router but that is not advertised via a PIO in the received RA, proceed as follows:
 - * Increment the LTA counter by one.
 - * IF LTA >= LAT_RAS_INVALID, then:
 - + IF this is the only router advertising this prefix, set the "Valid Lifetime" of this prefix to 0. This will cause the removal of all addresses for this prefix and of any routes for this prefix associated with the this router.
 - + ELSE IF this prefix has been advertised by multiple neighboring routers, simply disassociate this prefix with

this particular router. This will cause the fate of this prefix to depend on the other routers.

- * ELSE IF `LTA >= LTA_RAS_UNPREFER`, then:
 - + IF this is the only router advertising this prefix, set the "Preferred Lifetime" of this prefix to 0. This will cause the corresponding addresses to become un-preferred/deprecated.
 - + ELSE IF this prefix has been advertised by multiple neighboring routers, simply disassociate this prefix with this particular router. This will cause the fate of this prefix to depend on the other routers.

[Appendix A](#) illustrates a possible packet exchange and the operation of the algorithm for a typical scenario.

NOTES:

- o The processing of RAs that do not contain any PIOs with the "A" bit set remains unaffected.
- o The aforementioned processing assumes that PIOs will be spread among **at most** `LTA_RAS_UNPREFER` RAs.

We know of no implementations that intentionally spread network configuration information among multiple RAs, and know of no real deployment scenarios where the amount of SLAAC information requires that it be split into multiple Router Advertisement Messages. Therefore, while we have been conservative in selecting the default values for `LTA_RAS_UNPREFER` and `LTA_RAS_INVALID` as 2 and 4 (respectively), we believe that all real-world scenarios would be able to set them to 1 and 2 (respectively), without any impact on robustness. [Section 4.4](#) provides additional clarifications regarding how information is conveyed in Router Advertisement messages.

- o If the only prefix that has so far been advertised on the local network is the prefix that has become stale, and there is no other prefix being advertised, the traditional processing is unaffected (the mechanism discussed in this document will **never** be triggered because received RAs will not contain other PIOs with the "A" bit set). The rationale here is that it is better to have some address, than no address at all.
- o The specified modification takes the conservative approach of first setting the "Preferred Lifetime" to 0 (such that addresses

become non-preferred), and eventually setting the "Valid Lifetime" to 0 (such that addresses are completely removed). Once the addresses for this prefix have been removed, associated routes incorporated by the original RA messages SHOULD also be removed.

- o In cases where this scenario has been triggered by a CPE router crashing and rebooting, it would take hosts less than one minute to mark the corresponding addresses as "not preferred" (when using the default value for LTA_RAS_UNPREFER), and less than five minutes to completely remove such addresses from the system (when using the default value for LAT_RAS_INVALID).

[Section 6.2.4 of \[RFC4861\]](#) specifies that when an interface becomes an advertising interface, the first few unsolicited RAs (up to MAX_INITIAL_RTR_ADVERTISEMENTS, specified as 3) will be sent at intervals of at most MAX_INITIAL_RTR_ADVERT_INTERVAL (specified as 16 seconds). This means that, using the default value for LTA_RAS_UNPREFER (LTA_RAS_UNPREFER_DEFAULT=2), in the worst-case scenario it would take hosts 32 seconds to mark stale addresses as "not preferred". The fourth unsolicited RA will be sent after a random amount of time between MinRtrAdvInterval (default: $0.33 * \text{MaxRtrAdvInterval}$) and MaxRtrAdvInterval (default: 600 seconds) has elapsed, thus meaning that, when using the default value for LTA_RAS_INVALID (LTA_RAS_INVALID_DEFAULT=4) the stale addresses would be removed after between 3.3 and 10 minutes of being marked as "not preferred".

5. IANA Considerations

This document has no actions for IANA.

6. Security Considerations

An attacker that could impersonate a router could forge multiple RA packets that contain PIOs for prefixes that are currently not advertised on the local network and that fail to include previously-advertised prefixes, to trigger the mechanism specified in this document (and thus cause existing addresses to be deprecated, and eventually removed). However, an attacker that can impersonate a router could more easily deprecate addresses by advertising the legitimate prefixes with the "Preferred Lifetime" set to 0, or perform a plethora of other possible Denial of Service attacks based on forged RA messages. Therefore, when attacks based on forged RA packets are a concern, technologies such as RA-Guard [\[RFC6105\]](#) [\[RFC7113\]](#) should be deployed.

Capping the "Valid Lifetime" and "Preferred Lifetime" at hosts may help limit the duration of the effects of non-sustained attacks that employ forged RAs with PIOs, since hosts would now recover in a timelier manner.

7. Acknowledgments

The authors would like to thank (in alphabetical order) Mikael Abrahamsson, Luis Balbinot, Brian Carpenter, Owen DeLong, Gert Doering, Nick Hilliard, Bob Hinden, Philip Homburg, Lee Howard, Christian Huitema, Jen Linkova, Albert Manfredi, Jordi Palet Martinez, Michael Richardson, Mark Smith, Tarko Tikan, and Ole Troan, 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.

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.

The problem discussed in this document has been previously documented by Jen Linkova in [[I-D.linkova-6man-default-addr-selection-update](#)], and also in [[RIPE-690](#)].

8. References

8.1. 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>>.
- [RFC8028] Baker, F. and B. Carpenter, "First-Hop Router Selection by Hosts in a Multi-Prefix Network", [RFC 8028](#), DOI 10.17487/RFC8028, November 2016, <<https://www.rfc-editor.org/info/rfc8028>>.

- [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>>.
- [RFC8504] Chown, T., Loughney, J., and T. Winters, "IPv6 Node Requirements", [BCP 220](#), [RFC 8504](#), DOI 10.17487/RFC8504, January 2019, <<https://www.rfc-editor.org/info/rfc8504>>.

8.2. Informative References

- [I-D.ietf-v6ops-cpe-slaac-renum]
Gont, F., Zorz, J., and R. Patterson, "Improving the Reaction of Customer Edge Routers to Renumbering Events", [draft-ietf-v6ops-cpe-slaac-renum-00](#) (work in progress), February 2020.
- [I-D.ietf-v6ops-slaac-renum]
Gont, F., Zorz, J., and R. Patterson, "Reaction of Stateless Address Autoconfiguration (SLAAC) to Flash-Renumbering Events", [draft-ietf-v6ops-slaac-renum-00](#) (work in progress), March 2020.
- [I-D.linkova-6man-default-addr-selection-update]
Linkova, J., "Default Address Selection and Subnet Renumbering", [draft-linkova-6man-default-addr-selection-update-00](#) (work in progress), March 2017.
- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", [BCP 38](#), [RFC 2827](#), DOI 10.17487/RFC2827, May 2000, <<https://www.rfc-editor.org/info/rfc2827>>.
- [RFC5927] Gont, F., "ICMP Attacks against TCP", [RFC 5927](#), DOI 10.17487/RFC5927, July 2010, <<https://www.rfc-editor.org/info/rfc5927>>.
- [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>>.
- [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>>.

[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>>.

[RIPE-690] Zorz, J., Zorz, S., Drazumeric, P., Townsley, M., Alston, J., Doering, G., Palet, J., Linkova, J., Balbinot, L., Meynell, K., and L. Howard, "Best Current Operational Practice for Operators: IPv6 prefix assignment for end-users - persistent vs non-persistent, and what size to choose", RIPE 690, October 2017, <<https://www.ripe.net/publications/docs/ripe-690>>.

Appendix A. Sample Timeline for Host Processing of RAs

This section shows a sample packet exchange that illustrates the algorithm specified in [Section 4](#):


```

Router                                     Host
  RA, PIO={2001:DB8:1::/64, L=1, A=1}
----->
                                     [Host configures addrs
                                     for this prefix]
                                     LTA=0

  RA, PIO={2001:DB8:1::/64, L=1, A=1}
----->
                                     [Normal proc. of RA]
                                     .
                                     .

[Router reboots]

  RA, PIO={2001:DB8:2::/64, L=1, A=1}
-----> {2001:DB8:1::/64,
                                     LTA=1}
                                     .
                                     .
  RA, PIO={2001:DB8:2::/64, L=1, A=1}
-----> {2001:DB8:1::/64,
                                     LTA=2}
                                     LTA==LTA_RAS_UNPREFER
                                     Pref. Lftime=0
                                     .
                                     .
  RA, PIO={2001:DB8:2::/64, L=1, A=1}
-----> {2001:DB8:1::/64,
                                     LTA=3}
                                     .
                                     .
  RA, PIO={2001:DB8:2::/64, L=1, A=1}
-----> {2001:DB8:1::/64,
                                     LTA=4}
                                     LTA==LTA_RAS_INVALID
                                     Valid Lftime=0
                                     (Addr. Removed!)

```

[Appendix B](#). Analysis of Some Suggested Workarounds

[This section is to be removed before publication of this document as an RFC].

During the discussion of this document, some alternative workarounds were suggested on the 6man mailing-list. The following subsections analyze these suggested workarounds, in the hopes of avoiding rehashing the same discussions.

B.1. On a Possible Reaction to ICMPv6 Error Messages

It has been suggested that if configured addresses become stale, a CPE enforcing ingress/egress filtering ([BCP38](#)) ([\[RFC2827\]](#)) could send ICMPv6 Type 1 (Destination Unreachable) Code 5 (Source address failed ingress/egress policy) error messages to the sending node, and that, upon receipt of such error messages, the sending node could perform heuristics that might help to mitigate the problem discussed in this document.

The aforementioned proposal has a number of drawbacks and limitations:

- o It assumes that the CPE routers enforce ingress/egress filtering [\[RFC2827\]](#). While this is desirable behaviour, it cannot be relied upon.
- o It assumes that if the CPE enforces ingress/egress filtering, the CPE will signal the packet drops to the sending node with ICMPv6 Type 1 (Destination Unreachable) Code 5 (Source address failed ingress/egress policy) error messages. While this may be desirable, [\[RFC2827\]](#) does not suggest signaling the packet drops with ICMPv6 error messages, let alone the use of specific error messages (such as Type 1 Code 5) as suggested.
- o ICMPv6 Type 1 Code 5 could be interpreted as the employed address being stale, but also as a selected route being inappropriate/suboptimal. If the later, un-preferring addresses or removing addresses upon receipt of these error messages would be inappropriate.
- o Reacting to these error messages would create a new attack vector that could be exploited from remote networks. This is of particular concern since ICMP-based attacks do not even require that the Source Address of the attack packets be spoofed [\[RFC5927\]](#).

B.2. On a Possible Improvement to Source Address Selection

[\[RFC6724\]](#) specifies source address selection (SAS) for IPv6. Conceptually, it sorts the candidate set of source addresses for a given destination, based on a number of pair-wise comparison rules that must be successively applied until there is a "winning" address.

An implementation might improve source address selection, and prefer the most-recently advertised information. In order to incorporate the "freshness" of information in source address selection, an implementation would be updated as follows:

- o The node is assumed to maintain a timer/counter that is updated at least once per second. For example, the `time(2)` function from unix-like systems could be employed for this purpose.
- o The local information associated with each prefix advertised via RAs on the local network is augmented with a "LastAdvertised" timestamp value. Whenever an RA with a PIO with the "A" bit set for such prefix is received, the "LastAdvertised" timestamp is updated with the current value of the timer/counter.
- o [[RFC6724](#)] is updated such that this rule is incorporated:

Rule 7.5: Prefer fresh information If one of the two source addresses corresponds to a prefix that has been more recently advertised, say `LastAdvertised(SA) > LastAdvertised(SA)`, then prefer that address (SA in our case).

A clear benefit of this approach is that a host will normally prefer "fresh" addresses over possibly stale addresses.

However, there are a number of drawbacks associated with this approach:

- o In scenarios where multiple prefixes are being advertised on the same LAN segment, the new SAS rule is **guaranteed** to result in non-deterministic behaviour, with hosts frequently changing the default source address. This is certainly not desirable from a troubleshooting perspective.
- o Since the rule must be incorporated before "Rule 8: Use longest matching prefix" from [[RFC6724](#)], it may lead to suboptimal paths.
- o This new rule may help to improve the selection of a source address, but it does not help with the housekeeping (garbage collection) of configured information:
 - * If the stale prefix is re-used in another network, nodes employing stale addresses and routes for this prefix will be unable to communicate with the new "owner" of the prefix, since the stale prefix will most likely be considered "on-link".
 - * Given that the currently recommended default value for the "Valid Lifetime" of PIOs is 2592000 seconds (30 days), it would take too long for hosts to remove the configured addresses and routes for the stale prefix. While the proposed update in [Section 4.1](#) of this document would mitigate this problem, the lifetimes advertised by the local SLAAC router are not under the control of hosts.

As a result, updating IPv6 source address selection does not relieve nodes from improving their SLAAC implementations as specified in [Section 4](#), if at all desirable. On the other hand, the algorithm specified in [Section 4.5](#) would result in Rule 3 of [[RFC6724](#)] employing fresh addresses, without leading to non-deterministic behaviour.

Authors' Addresses

Fernando Gont
SI6 Networks / UTN-FRH
Seguro y Habana 4310, 7mo Piso
Villa Devoto, Ciudad Autonoma de Buenos Aires
Argentina

Email: fgont@si6networks.com
URI: <https://www.si6networks.com>

Jan Zorz
Go6 Institute
Frankovo naselje 165
Skofja Loka 4220
Slovenia

Email: jan@go6.si
URI: <https://www.go6.si>

Richard Patterson
Sky UK

Email: richard.patterson@sky.uk

