

IPv6 Operations Working Group (v6ops)
Internet-Draft
Intended status: Informational
Expires: April 1, 2021

F. Gont
SI6 Networks
J. Zorz
6connect
R. Patterson
Sky UK
September 28, 2020

**Reaction of Stateless Address Autoconfiguration (SLAAC) to Flash-
Renumbering Events
draft-ietf-v6ops-slaac-renum-04**

Abstract

In scenarios where network configuration information related to IPv6 prefixes becomes invalid without any explicit signaling of that condition (such as when a CPE crashes and reboots without knowledge of the previously-employed prefixes), nodes on the local network will continue using stale prefixes for an unacceptably long time (on the order of several days), thus resulting in connectivity problems. This document documents this issue and discusses operational workarounds that may help to improve network robustness. Additionally, it highlights areas where further work may be needed.

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 April 1, 2021.

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.	Analysis of the Problem	5
2.1.	Use of Dynamic Prefixes	5
2.2.	Default Timer Values in IPv6 Stateless Address Autoconfiguration (SLAAC)	6
2.3.	Recovering from Stale Network Configuration Information .	6
2.4.	Lack of Explicit Signaling about Stale Information . . .	6
2.5.	Interaction Between DHCPv6-PD and SLAAC	7
3.	Operational Mitigations	7
3.1.	Stable Prefixes	7
3.2.	SLAAC Parameter Tweaking	8
4.	Future Work	9
5.	IANA Considerations	9
6.	Security Considerations	9
7.	Acknowledgments	9
8.	References	10
8.1.	Normative References	10
8.2.	Informative References	11
	Authors' Addresses	12

[1.](#) Introduction

IPv6 Stateless address autoconfiguration (SLAAC) [[RFC4862](#)] conveys information about prefixes to be employed for address configuration via Prefix Information Options (PIOs) sent in Router Advertisement (RA) messages. IPv6 largely assumes prefix stability, with network renumbering only taking place in a planned manner, with old/stale prefixes being phased-out via reduced prefix lifetimes, and new prefixes (with longer lifetimes) being introduced at the same time. However, there are several scenarios that may lead to the so-called "flash-renumbering" events, where the prefix employed by a network suddenly becomes invalid and replaced by a new prefix. In some of these scenarios, the local router producing the network renumbering event may try to deprecate the currently-employed prefixes (by explicitly signaling the network about the renumbering event), whereas in other scenarios it may be unable to do so.

In scenarios where network configuration information related to IPv6 prefixes becomes invalid without any explicit signaling of that condition, nodes on the local network will continue using stale prefixes for an unacceptably long period of time, thus resulting in connectivity problems.

Scenarios where this problem may arise include, but are not limited to, the following:

- o The most common IPv6 deployment scenario for residential or small office networks is that in which a CPE router employs DHCPv6 Prefix Delegation (DHCPv6-PD) [[RFC8415](#)] to request a prefix from an Internet Service Provider (ISP), and a sub-prefix of the leased prefix is advertised on the LAN-side of the CPE router via Stateless Address Autoconfiguration (SLAAC) [[RFC4862](#)]. In scenarios where the CPE router crashes and reboots, the CPE may obtain (via DHCPv6-PD) a different prefix from the one previously leased, and therefore advertise (via SLAAC) the new prefix on the LAN side. Hosts will typically configure addresses for the new prefix, but will normally retain and may actively employ the addresses configured for the previously-advertised prefix, since their associated Preferred Lifetime and Valid Lifetime allow them to do so.
- o A router (e.g. Customer Edge router) may advertise autoconfiguration prefixes corresponding to prefixes learned via DHCPv6-PD with constant PIO lifetimes that are not synchronized with the DHCPv6-PD lease time (as required in [Section 6.3 of \[RFC8415\]](#)). While this behavior violates the aforementioned requirement from [[RFC8415](#)], it is not an unusual behavior, particularly when e.g. DHCPv6-PD is implemented in a different software module than the SLAAC router component.
- o A switch-port the host is connected to may be moved to another subnet (VLAN) as a result of manual switch-port reconfiguration or 802.1x re-authentication. In particular, there has been evidence that some 802.1x supplicants do not reset network settings after successful 802.1x authentication. So if a host had failed 802.1x authentication for some reason, was placed in a "quarantine" VLAN and then got successfully authenticated later on, it might end up having IPv6 addresses from both old ("quarantine") and new VLANs.
- o During the planned network renumbering, a router may be configured to send an RA with the Preferred Lifetime for the "old" Prefix Information Option (PIO) set to zero and the new PIO having non-zero Preferred Lifetime. However, due to unsolicited RAs being sent as all-hosts multicast and multicast being rather unreliable

on busy wifi networks, the RA may not be received by a host (or set of hosts).

- o Automated device config management system performs periodic config pushes to network devices. In these scenarios, network devices may simply immediately forget their previous configuration, rather than withdrawing it gracefully. If such a push results in changing the subnet configured on a particular network, hosts attached to that network would not get notified about the subnet change, and their addresses from the "old" prefix will not be deprecated. A related scenario is the incorrect network renumbering where a network administrator renumbers a network by simply removing the "old" prefix from the configuration and configuring a new prefix instead.

Lacking any explicit signaling to deprecate the previously-advertised prefixes, hosts may continue to employ the previously-configured addresses, which will typically result in packets being blackholed -- whether because of egress-filtering by the CPE or ISP, or packets being discarded or routed elsewhere.

We note that the default values for the "Preferred Lifetime" and "Valid Lifetime" of PIOs, as specified in [[RFC4861](#)], are:

- o Preferred Lifetime (AdvPreferredLifetime): 604800 seconds (7 days)
- o Valid Lifetime (AdvValidLifetime): 2592000 seconds (30 days)

This means that in the aforementioned scenarios, the stale addresses would be retained and also actively employed for new communications instances for an unacceptably long period of time (one month, and one week, respectively), leading to interoperability problems, instead of hosts transitioning to the newly-advertised prefix(es) in a timelier manner.

Some devices have implemented ad-hoc mechanisms to address this problem, such as sending RAs to deprecate apparently-stale prefixes when the device receives any packets employing a source address from a prefix not currently advertised for address configuration on the local network [[FRITZ](#)]. However, this may introduce other interoperability problems, particularly in multihomed/multiprefix scenarios. This is a clear indication that advice in this area is warranted.

Unresponsiveness to these "flash-renumbering" events is caused by the inability of the network to deprecate stale information, as well as by the inability of hosts to react to network configuration changes in a timelier manner. Clearly, it would be desirable that these

flash-renumbering scenarios do not occur, and that, when they do occur, that hosts are explicitly notified of their occurrence. However, for robustness reasons, it is paramount for hosts to be able to recover from stale configuration information even when these flash-renumbering events occur and the network is unable to explicitly notify hosts about such conditions.

[Section 2](#) analyzes this problem in more detail. [Section 3](#) describes possible operational mitigations. [Section 4](#) describes possible future work to mitigate the aforementioned problem.

2. Analysis of the Problem

As noted in [Section 1](#), the problems discussed in this document are exacerbated by the default values of some protocol parameters and other factors. The following sections analyze each of them in detail.

2.1. Use of Dynamic Prefixes

In the residential or small office scenario, the problem discussed in this document would be avoided if DHCPv6-PD would lease "stable" prefixes. However, a recent survey [[UK-NOF](#)] indicates that 37% of the responding ISPs employ dynamic prefixes. That is, dynamic IPv6 prefixes are an operational reality.

Deployment reality aside, there are a number of possible issues associated with stable prefixes:

- o Provisioning systems may be unable to deliver stable IPv6 prefixes.
- o While there is a range of information that may be employed to correlate network activity [[RFC7721](#)], the use of stable prefixes clearly simplifies network activity correlation, and may essentially render features such as "temporary addresses" [[RFC4941](#)] irrelevant.
- o There may be existing advice for ISPs to deliver dynamic IPv6 prefixes *by default* (see e.g. [[GERMAN-DP](#)]) over privacy concerns associated with stable prefixes.

For a number of reasons (such as the ones stated above), IPv6 deployments may employ dynamic prefixes (even at the expense of the issues discussed in this document), and that there might be scenarios in which the dynamics of a network are such that the network exhibits the behaviour of dynamic prefixes. Rather than trying to regulate

how operators may run their networks, this document aims at improving network robustness in the deployed Internet.

2.2. Default Timer Values in IPv6 Stateless Address Autoconfiguration (SLAAC)

IPv6 SLAAC employs the following default PIO lifetime values:

- o Preferred Lifetime (AdvPreferredLifetime): 604800 seconds (7 days)
- o Valid Lifetime (AdvValidLifetime): 2592000 seconds (30 days)

Under problematic circumstances, such as where the corresponding network information has become stale without any explicit signal from the network (as described in [Section 1](#)), it will take a host 7 days (one week) to deprecate the corresponding addresses, and 30 days (one month) to eventually invalidate and remove any addresses configured for the stale prefix. This means that it will typically take hosts an unacceptably long period of time (on the order of several days) to recover from these scenarios.

2.3. Recovering from Stale Network Configuration Information

SLAAC hosts are unable to recover from stale network configuration information for a number of reasons:

- o Item "e)" of [Section 5.5.3 of \[RFC4862\]](#) specifies that an RA may never reduce the "RemainingLifetime" to less than 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 Preferred Lifetime of an address can be reduced to any value to avoid using a stale prefix for new communications.
- o In the absence of explicit signalling from SLAAC routers (such as sending PIOs with a "Preferred Lifetime" set to 0), SLAAC hosts fail to recover from stale configuration information in a timely manner. However, when a network element is able to explicitly signal the renumbering event, it will only be able to deprecate the stale prefix, but not to invalidate the prefix in question. Therefore, communication with the new "owners" of the stale prefix will not be possible, since the stale prefix will still be considered "on-link".

2.4. Lack of Explicit Signaling about Stale Information

Whenever prefix information has changed, a SLAAC router should not only advertise the new information, but should also advertise the stale information with appropriate lifetime values (both "Preferred

Lifetime" and "Valid Lifetime" set to 0). This would provide explicit signaling to SLAAC hosts to remove the stale information (including configured addresses and routes). However, in scenarios such as when a CPE router crashes and reboots, the CPE router may have no knowledge about the previously-advertised prefixes, and thus may be unable to advertise them with appropriate lifetimes (in order to deprecate them).

However, we note that, as discussed in [Section 2.3](#), PIOs with small Valid Lifetimes will not lower the Valid Lifetime to any value shorter than two hours (as per [\[RFC4862\]](#)). Therefore, even if a SLAAC router tried to explicitly signal the network about the stale configuration information via RAs, implementations compliant with [\[RFC4862\]](#) would deprecate the corresponding prefixes, but would fail to invalidate them.

NOTE:

Some implementations have been updated to honor small PIO lifetimes values, as proposed in [\[I-D.ietf-6man-slaac-renum\]](#). For example, please see [\[Linux\]](#).

[2.5. Interaction Between DHCPv6-PD and SLAAC](#)

While DHCPv6-PD is normally employed along with SLAAC, the interaction between the two protocols is largely unspecified. Not unusually, the two protocols are implemented in two different software components with the interface between the two implemented by some sort of script that feeds the SLAAC implementation with values learned from DHCPv6-PD.

At times, the prefix lease time is fed as a constant value to the SLAAC router implementation, meaning that, eventually, the prefix lifetime advertised on the LAN side will span *past* the DHCPv6-PD lease time. This is clearly incorrect, since the SLAAC router implementation would be allowing the use of such prefixes for a longer time than it has been granted usage of those prefixes via DHCPv6-PD.

[3. Operational Mitigations](#)

The following subsections discuss possible operational workarounds for the aforementioned problems.

[3.1. Stable Prefixes](#)

As noted in [Section 2.1](#), the use of stable prefixes would eliminate the issue in *some* of the scenarios discussed in [Section 1](#) of this document, such as the typical home network deployment. However, even

in such scenarios, there might be reasons for which an administrator may want or may need to employ dynamic prefixes

3.2. SLAAC Parameter Tweaking

An operator may wish to override some SLAAC parameters such that, under normal circumstances (including some packet loss), the timers will be refreshed/reset, but in the presence of network faults (such as network configuration information becoming stale without explicit signaling), the timers go off and trigger some fault recovering action (such as deprecating the corresponding addresses and subsequently invalidating/removing them).

The following router configuration variables from [[RFC4861](#)] (corresponding to the "lifetime" parameters of PIOs) could be overridden as follows:

AdvPreferredLifetime: 2700 seconds (45 minutes)

AdvValidLifetime: 5400 seconds (90 minutes)

NOTES:

A CPE router advertising a sub-prefix of a prefixed leased via DHCPv6-PD will periodically refresh the Preferred Lifetime and the Valid Lifetime of an advertised prefix to AdvPreferredLifetime and AdvValidLifetime, respectively, as long as the resulting lifetime of the corresponding prefixes does not extend past the DHCPv6-PD lease time [[I-D.ietf-v6ops-cpe-slaac-renum](#)].

RATIONALE:

- * In the context of [[RFC8028](#)], where it is clear that use of addresses configured for a given prefix is tied to using the next-hop router that advertised the prefix, it does not make sense for the "Preferred Lifetime" of a PIO to be larger than the "Router Lifetime" (AdvDefaultLifetime) of the corresponding Router Advertisement messages. The "Valid Lifetime" is set to a much larger value to cope with transient network problems.
- * Lacking RAs that refresh information, addresses configured for advertised prefixes become deprecated in a timelier manner, and thus Rule 3 of [[RFC6724](#)] causes other configured addresses (if available) to be used instead.
- * We note that lowering the default values for the "Valid Lifetime" helps reduce the amount of time a host may maintain stale information and the amount of time an advertising router would need to advertise stale prefixes to deprecate them, while

reducing the default "Preferred Lifetime" would reduce the amount of time it takes for a host to prefer other working prefixes (see [Section 12 of \[RFC4861\]](#)). However, while the values suggested in this section are an improvement over the default values specified in [\[RFC4861\]](#), they represent a trade-off among a number of factors, including responsiveness, possible impact on the battery life of connected devices [\[RFC7772\]](#), etc. Thus, they may or may not provide sufficient mitigation to the problem discussed in this document.

4. Future Work

Improvement in Customer Edge Routers [\[RFC7084\]](#) such that they can signal the network about stale prefixes and deprecate them accordingly can help mitigate the problem discussed in this document for the "home network" scenario. Such work is currently being pursued in [\[I-D.ietf-v6ops-cpe-slaac-renum\]](#).

Improvements in the SLAAC protocol [\[RFC4862\]](#) and other algorithms such as "Default Address Selection for IPv6" [\[RFC6724\]](#) would help improve network robustness. Such work is currently being pursued in [\[I-D.ietf-6man-slaac-renum\]](#).

The aforementioned work is considered out of the scope of this present document, which only focuses on documenting the problem and discussing operational mitigations.

5. IANA Considerations

This document has no actions for IANA.

6. Security Considerations

This document discusses a problem that may arise in scenarios where flash-renumbering events occur, and proposes workarounds to mitigate the aforementioned problems. This document does not introduce any new security issues.

7. Acknowledgments

The authors would like to thank (in alphabetical order) Brian Carpenter, Owen DeLong, Guillermo Gont, Philip Homburg, Warren Kumari, Ted Lemon, Juergen Schoenwaelder, Klaas Wierenga, and Dale Worley, for providing valuable comments on earlier versions of this document.

The authors would like to thank (in alphabetical order) Mikael Abrahamsson, Luis Balbinot, Brian Carpenter, Tassos Chatzithomaoglou,

Uesley Correa, Owen DeLong, Gert Doering, Fernando Frediani, Steinar Haug, Nick Hilliard, Philip Homburg, Lee Howard, Christian Huitema, Ted Lemon, Albert Manfredi, Jordi Palet Martinez, Michael Richardson, Mark Smith, Tarko Tikan, and Ole Troan, for providing valuable comments on [[I-D.gont-6man-slaac-renum](#)], on which this document is based.

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 have 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](#)]. [Section 1](#) borrows text from [[I-D.linkova-6man-default-addr-selection-update](#)], authored by Jen Linkova.

8. References

8.1. Normative References

- [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>>.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", [RFC 4941](#), DOI 10.17487/RFC4941, September 2007, <<https://www.rfc-editor.org/info/rfc4941>>.
- [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>>.
- [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>>.

8.2. Informative References

- [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/>>.
- [GERMAN-DP] BFDI, "Einführung von IPv6 Hinweise für Provider im Privatkundengeschäft und Hersteller", Entschliessung der 84. Konferenz der Datenschutzbeauftragten des Bundes und der Länder am 7./8. November 2012 in Frankfurt (Oder), November 2012, <http://www.bfdi.bund.de/SharedDocs/Publikationen/Entschliessungssammlung/DSBundLaender/84DSK_EinfuehrungIPv6.pdf?__blob=publicationFile>.
- [I-D.gont-6man-slaac-renum] Gont, F., Zorz, J., and R. Patterson, "Improving the Robustness of Stateless Address Autoconfiguration (SLAAC) to Flash Renumbering Events", [draft-gont-6man-slaac-renum-08](#) (work in progress), May 2020.
- [I-D.ietf-6man-slaac-renum] Gont, F., Zorz, J., and R. Patterson, "Improving the Robustness of Stateless Address Autoconfiguration (SLAAC) to Flash Renumbering Events", [draft-ietf-6man-slaac-renum-01](#) (work in progress), August 2020.
- [I-D.ietf-v6ops-cpe-slaac-renum] Gont, F., Zorz, J., Patterson, R., and B. Volz, "Improving the Reaction of Customer Edge Routers to Renumbering Events", [draft-ietf-v6ops-cpe-slaac-renum-04](#) (work in progress), August 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.
- [Linux] Gont, F., "[net-next] ipv6: Honor all IPv6 PIO Valid Lifetime values", Post to the netdev mailing-list <http://vger.kernel.org/vger-lists.html>, April 2020, <<https://patchwork.ozlabs.org/project/netdev/patch/20200419122457.GA971@archlinux-current.localdomain/>>.

- [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>>.
- [RFC7084] Singh, H., Beebee, W., Donley, C., and B. Stark, "Basic Requirements for IPv6 Customer Edge Routers", [RFC 7084](#), DOI 10.17487/RFC7084, November 2013, <<https://www.rfc-editor.org/info/rfc7084>>.
- [RFC7721] Cooper, A., Gont, F., and D. Thaler, "Security and Privacy Considerations for IPv6 Address Generation Mechanisms", [RFC 7721](#), DOI 10.17487/RFC7721, March 2016, <<https://www.rfc-editor.org/info/rfc7721>>.
- [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>>.
- [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>>.
- [UK-NOF] Palet, J., "IPv6 Deployment Survey (Residential/Household Services) How IPv6 is being deployed?", UK NOF 39, January 2018, <<https://indico.uknof.org.uk/event/41/contributions/542/attachments/712/866/bcop-ipv6-prefix-v9.pdf>>.

Authors' Addresses

Fernando Gont
SI6 Networks
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
6connect

Email: jan@connect.com

Richard Patterson
Sky UK

Email: richard.patterson@sky.uk