

Internet Engineering Task Force (IETF)
Request for Comments: 9096
BCP: 234
Updates: 7084
Category: Best Current Practice
ISSN: 2070-1721

F. Gont
SI6 Networks
J. Å½orÅ¼
6connect
R. Patterson
Sky UK
B. Volz
Individual Contributor
August 2021

Improving the Reaction of Customer Edge Routers to IPv6 Renumbering Events

Abstract

This document specifies improvements to Customer Edge routers that help mitigate the problems that may arise when network configuration information becomes invalid without any explicit signaling of that condition to the local nodes. This document updates RFC 7084.

Status of This Memo

This memo documents an Internet Best Current Practice.

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

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9096>.

Copyright Notice

Copyright (c) 2021 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. Requirements Language
 3. Improved Customer Edge Router Behavior
 - 3.1. Automatic DHCPv6 RELEASES
 - 3.2. Stability of IAIDs
 - 3.3. Interface between the WAN Side and LAN Side
 - 3.4. LAN-Side Option Lifetimes
 - 3.5. Signaling Stale Configuration Information
 4. Recommended Option Lifetimes Configuration Values
 5. IANA Considerations
 6. Security Considerations
 7. References
 - 7.1. Normative References
 - 7.2. Informative References
- Acknowledgments

1. Introduction

In scenarios where network configuration information becomes invalid without any explicit signaling of that condition (such as when a Customer Edge (CE) router crashes and reboots without knowledge of the previously employed configuration information), hosts on the local network will continue using stale information for an unacceptably long period of time, thus resulting in connectivity problems. This problem is documented in detail in [RFC8978].

This document specifies improvements to CE routers that help mitigate the aforementioned problem for residential and small office scenarios. It specifies recommendations for the default behavior of CE routers but does not preclude the availability of configuration knobs that might allow an operator or user to manually configure the CE router to deviate from these recommendations. This document updates RFC 7084.

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. Improved Customer Edge Router Behavior

This section specifies and clarifies requirements for CE routers that can help mitigate the problem discussed in Section 1, particularly when they employ prefixes learned via DHCPv6 Prefix Delegation (DHCPv6-PD) [RFC8415] on the WAN side with Stateless Address Autoconfiguration (SLAAC) [RFC4862] or DHCPv6 [RFC8415] on the LAN side. The recommendations in this document help improve robustness at the CE router (on which the user or ISP may have no control) and do not preclude implementation of host-side improvements such as those specified in [6MAN-SLAAC-RENUM].

This document specifies additional WAN-side prefix-delegation (WPD) requirements to those specified in [RFC7084]:

WPD-9: CE routers SHOULD NOT automatically send DHCPv6-PD RELEASE messages upon restart events. See Section 3.1 for further details.

WPD-10: CE routers MUST by default use a WAN-side Identity Association Identifier (IAID) value that is stable between CE router restarts, DHCPv6 client restarts, or interface state changes (e.g., transient PPP interfaces), unless the CE router employs the IAID techniques discussed in Section 4.5 of [RFC7844]. See Section 3.2 for further details.

This document also replaces LAN-side requirement L-13 from [RFC7084] with:

L-13: CE routers MUST signal stale configuration information as specified in Section 3.5.

Finally, this document specifies the following additional LAN-side requirements to those from [RFC7084]:

L-15: CE routers MUST NOT advertise prefixes via SLAAC or assign addresses or delegate prefixes via DHCPv6 on the LAN side using lifetimes that exceed the remaining lifetimes of the corresponding prefixes learned on the WAN side via DHCPv6-PD. For more details, see Section 3.3.

L-16: CE routers SHOULD advertise capped SLAAC option lifetimes, capped DHCPv6 IA Address option lifetimes, and capped IA Prefix

option lifetimes, as specified in Section 3.4.

3.1. Automatic DHCPv6 RELEASES

Some CE routers are known to automatically send DHCPv6-PD RELEASE messages upon restart events. However, this may inadvertently trigger a flash-renumbering scenario, along with the associated problems discussed in [RFC8978], that this document attempts to mitigate.

As a result, requirement WPD-9 from Section 3 specifies that CE routers SHOULD NOT automatically send DHCPv6-PD RELEASE messages upon restart events.

3.2. Stability of IAIDs

[RFC8415] requires that the IAID for an IA MUST be consistent across restarts of the DHCP client. However, some popular CE routers are known to select new random IAIDs, e.g., every time the underlying PPP session is established or when the device is rebooted. This could be the result of extrapolating the behavior described in [RFC7844] or simply a consequence of not storing IAIDs on stable storage along with failure to employ an algorithm that consistently generates the same IAID upon reboots. Thus, requirement WPD-10 from Section 3 prevents CE routers from inadvertently triggering flash-renumbering events on the local network.

3.3. Interface between the WAN Side and LAN Side

The "Preferred Lifetime" and "Valid Lifetime" of Prefix Information Options (PIOs) [RFC4861] corresponding to prefixes learned via DHCPv6-PD on the WAN side MUST NOT span past the remaining preferred and valid lifetimes of the corresponding DHCPv6-PD prefixes. This means that the "Preferred Lifetime" and the "Valid Lifetime" advertised in PIOs by the CE router MUST be dynamically adjusted such that they never span past the remaining preferred and valid lifetimes of the corresponding prefixes delegated via DHCPv6-PD on the WAN side.

Similarly, the "preferred-lifetime" and "valid-lifetime" of DHCPv6 IA Address options and DHCPv6 IA Prefix options employed with DHCPv6 on the LAN side MUST NOT span past the remaining preferred and valid lifetimes of the corresponding prefixes learned via DHCPv6-PD on the WAN side. This means that the "preferred-lifetime" and "valid-lifetime" of DHCPv6 IA Address options and DHCPv6 IA Prefix options employed with DHCPv6 on the LAN side MUST be dynamically adjusted such that they never span past the remaining preferred and valid lifetimes of the corresponding prefixes delegated to the CE router on the WAN side via DHCPv6-PD.

RATIONALE:

- * The lifetime values employed for the "Preferred Lifetime" (AdvPreferredLifetime) and "Valid Lifetime" (AdvValidLifetime) of SLAAC Prefix Information Options must never be larger than the remaining lifetimes of the corresponding prefixes (as learned via DHCPv6-PD on the WAN side). This is in line with the requirement from Section 6.3 of [RFC8415], which states:

In particular, if the delegated prefix or a prefix derived from it is advertised for stateless address autoconfiguration [RFC4862], the advertised preferred and valid lifetimes MUST NOT exceed the corresponding remaining lifetimes of the delegated prefix.

- * The lifetime values of prefixes advertised on the LAN side via SLAAC must be dynamically updated (rather than static values); otherwise, the advertised lifetimes would eventually span past the DHCPv6-PD lifetimes.
- * The same considerations apply for the "valid-lifetime" and "preferred-lifetime" of IA Address options and IA Prefix options

employed with DHCPv6 on the LAN side.

3.4. LAN-Side Option Lifetimes

CE routers SHOULD override the default lifetime values of Neighbor Discovery options that depend in any way on changes in the prefix employed for address configuration on the LAN side, and employ shorter lifetime values to improve the robustness to renumbering events, while complying with the requirements from Section 3.3 of this document and the recommendations in [RFC7772].

CE routers SHOULD set the "Router Lifetime" of Router Advertisement (RA) messages to ND_PREFERRED_LIMIT.

CE routers SHOULD also set the PIO "Preferred Lifetime" to the lesser of the remaining preferred lifetime of the corresponding prefix (see Section 3.3) and ND_PREFERRED_LIMIT, and set the PIO "Valid Lifetime" to the lesser of the remaining valid lifetime of the corresponding prefix and ND_VALID_LIMIT. Additionally, the "Route Lifetime" of Route Information Options (RIOs) [RFC4191], the "Lifetime" of Recursive DNS Server (RDNS) options [RFC8106], and the "Lifetime" of DNS Search List (DNSSL) options [RFC8106] SHOULD be set to the lesser of the longest remaining valid lifetime of a prefix (leased via DHCPv6 on the WAN side) and ND_VALID_LIMIT, if any of these options are included in Router Advertisement messages.

NOTE: In scenarios where the valid lifetime and the preferred lifetime of prefixes learned via DHCPv6 on the WAN side are always larger than ND_VALID_LIMIT and ND_PREFERRED_LIMIT, respectively, the lifetime values advertised on the LAN side will not experience actual changes.

The above text refers to the Neighbor Discovery options that are typically employed by CE routers. A CE router may need to apply the same policy for setting the lifetime of other Neighbor Discovery options it employs, if and where applicable.

CE routers providing stateful address configuration via DHCPv6 SHOULD set the "preferred-lifetime" of a DHCPv6 IA Address option to the lesser of the remaining preferred lifetime of the corresponding prefix (see Section 3.3) and ND_PREFERRED_LIMIT, and set the "valid-lifetime" of the same option to the lesser of the remaining valid lifetime of the corresponding prefix and ND_VALID_LIMIT.

CE routers providing DHCPv6-PD on the LAN side SHOULD set the "preferred-lifetime" of a DHCPv6 IA Prefix option to the lesser of the remaining preferred lifetime of the corresponding prefix (see Section 3.3) and ND_PREFERRED_LIMIT, and set the "valid-lifetime" of the same option to the lesser of the remaining valid lifetime of the corresponding prefix and ND_VALID_LIMIT.

RATIONALE:

- * The "Valid Lifetime" and "Preferred Lifetime" of PIOs have a direct impact on three different aspects:
 - The amount of time hosts may end up employing stale network configuration information (see [RFC8978]).
 - The amount of time CE routers need to persist trying to deprecate stale network configuration information (e.g., to handle cases where hosts miss Router Advertisement messages and thus still consider the stale information as valid).
 - The amount of information that CE routers need to maintain when, e.g., multiple crash-and-reboot events occur in the time span represented by the option lifetimes employed on the LAN side.
- * CE routers need not employ the (possibly long) WAN-side DHCPv6-PD lifetimes for the "Valid Lifetime" and "Preferred Lifetime" of

PIOs sent in Router Advertisement messages to advertise sub-prefixes of the leased prefix. Instead, CE routers SHOULD use shorter values for the "Valid Lifetime" and "Preferred Lifetime" of PIOs, since subsequent Router Advertisement messages will nevertheless refresh the associated lifetimes, leading to the same effective lifetimes as specified by the WAN-side DHCPv6-PD lifetimes.

- * Similarly, CE routers need not employ the (possibly long) WAN-side DHCPv6-PD lifetimes for the "valid-lifetime" and "preferred-lifetime" of IA Address options and IA Prefix options employed by DHCPv6 on the LAN side, since the renewal of bindings by DHCPv6 clients will lead to the same effective lifetimes as specified by the WAN-side DHCPv6-PD lifetimes.

3.5. Signaling Stale Configuration Information

When a CE router provides LAN-side address-configuration information via SLAAC:

- * A CE router sending RAs that advertise prefixes belonging to a dynamically learned prefix (e.g., via DHCPv6-PD) SHOULD record, on stable storage, the list of prefixes being advertised via PIOs on each network segment and the state of the "A" and "L" flags of the corresponding PIOs.
- * Upon changes to the advertised prefixes, and after bootstrapping, the CE router advertising prefix information via SLAAC proceeds as follows:
 - Any prefixes that were previously advertised by the CE router via PIOs in RA messages, but that have now become stale, MUST be advertised with PIOs that have the "Valid Lifetime" and the "Preferred Lifetime" set to 0 and the "A" and "L" bits unchanged.
 - The aforementioned advertisements MUST be performed for at least the "Valid Lifetime" previously employed for such prefixes. The CE router MUST advertise this information with unsolicited Router Advertisement messages, as described in Section 6.2.4 of [RFC4861], and MAY advertise this information via unicast Router Advertisement messages when possible and applicable.

NOTE: If requirement L-16 (Section 3) is followed, the "Valid Lifetime" need not be saved, and the stale prefix can simply be advertised for a period of ND_VALID_LIMIT.

- * CE routers receiving DHCPv6 IA Prefix options with a 0 "valid-lifetime" MUST advertise the corresponding sub-prefixes (as they would be generated for the same leased prefix with a non-zero lifetime) with PIOs with both the "Preferred Lifetime" and the "Valid Lifetime" set to 0, for at least the WAN-side DHCPv6-PD "valid-lifetime", or for a period of ND_VALID_LIMIT if the recommended lifetimes from Section 3.4 are employed.

When a CE router provides LAN-side DHCPv6 (address assignment or prefix delegation), then:

- * The CE router SHOULD record, on stable storage, the DHCPv6 address and delegated-prefix bindings corresponding to the LAN side.
- * If the CE router finds that the prefix to be employed for address assignment and/or prefix delegation has changed (e.g., upon a crash-and-reboot event) or the CE router receives DHCPv6 IA Prefix options with 0 lifetimes, the CE router MUST:
 - In Replies to DHCPv6 Request, Renew, and Rebind messages, send IA Address options or IA Prefix options (as appropriate) for any address assignments or prefix delegations for the stale prefixes. The aforementioned options MUST be sent with both

the "valid-lifetime" and the "preferred-lifetime" set to 0, for at least the "valid-lifetime" originally employed for them, or for a period of ND_VALID_LIMIT if the recommended lifetimes from Section 3.4 are employed.

- Initiate sending Reconfigure messages, if possible (i.e., client requests Reconfigure support and the CE router offers it), to those clients with address assignments or prefix delegations for the stale prefixes.

RATIONALE:

- * 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, a number of scenarios may lead to the so-called "flash-renumbering" events, where a prefix being employed on a network suddenly becomes invalid and replaced by a new prefix [RFC8978]. One such scenario is when an Internet Service Provider (ISP) employs dynamic prefixes and the CE router crashes and reboots. The requirements in this section are meant to allow CE routers to deprecate stale information in such scenarios.
- * The recommendations in this section expand from requirement L-13 in Section 4.3 of [RFC7084] and Section 6.3 of [RFC8415].
- * Hosts configuring addresses via SLAAC on the local network may employ addresses configured for the previously advertised prefixes for at most the "Valid Lifetime" of the corresponding PIOs of the last received Router Advertisement messages. Since Router Advertisement messages may be lost or fail to be received for various reasons, CE routers need to try to deprecate stale prefixes for a period of time equal to the "Valid Lifetime" of the PIO employed when originally advertising the prefix.
- * The requirements in this section to store information on stable storage are conveyed as "SHOULD" (as opposed to "MUST"), since they may represent a challenge for some implementations.
- * Advertising DHCPv6-leased prefixes with zero lifetimes on the LAN side would handle the case where a CE router has no stable storage but receives the prefixes via DHCPv6 with 0 lifetimes.
- * The above text does not include DHCPv6 Advertise messages sent in response to DHCPv6 Solicit messages, since Section 18.3.9 of [RFC8415] requires that a DHCPv6 server that is not going to assign an address or delegated prefix received as a hint in the Solicit message MUST NOT include that address or delegated prefix in the Advertise message. Additionally, any subsequent Request messages will trigger the response specified in this section and therefore cause the address or prefix to be deprecated.

4. Recommended Option Lifetimes Configuration Values

- * 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.

5. IANA Considerations

This document has no IANA actions.

6. Security Considerations

This document discusses a problem that may arise, e.g., in scenarios where dynamic IPv6 prefixes are employed, and it proposes improvements to CE routers [RFC7084] to mitigate the problem for residential or small office scenarios. It does not introduce new security issues; thus, the same security considerations as for [RFC4861], [RFC4862], [RFC7084], and [RFC8415] apply.

7. References

7.1. 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>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC7844] Huitema, C., Mrugalski, T., and S. Krishnan, "Anonymity Profiles for DHCP Clients", RFC 7844, DOI 10.17487/RFC7844, May 2016, <<https://www.rfc-editor.org/info/rfc7844>>.
- [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>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8415] Mrugalski, T., Siodelski, M., Volz, B., Yourtchenko, A., Richardson, M., Jiang, S., Lemon, T., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 8415, DOI 10.17487/RFC8415, November 2018, <<https://www.rfc-editor.org/info/rfc8415>>.

7.2. Informative References

- [6MAN-SLAAC-RENUM] Gont, F., Zorz, J., and R. Patterson, "Improving the Robustness of Stateless Address Autoconfiguration (SLAAC) to Flash Renumbering Events", Work in Progress, Internet-Draft, draft-ietf-6man-slaac-renum-02, 19 January 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-6man-slaac-renum-02>>.

- [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>>.
- [RFC8978] Gont, F., Å½orÅ¾, J., and R. Patterson, "Reaction of IPv6 Stateless Address Autoconfiguration (SLAAC) to Flash-Resubscribing Events", RFC 8978, DOI 10.17487/RFC8978, March 2021, <<https://www.rfc-editor.org/info/rfc8978>>.

Acknowledgments

The authors would like to thank Owen DeLong, Philip Homburg, Erik Kline, and Ted Lemon for their valuable help in improving this document via successive detailed reviews.

The authors would like to thank Mikael Abrahamsson, Luis Balbinot, Brian Carpenter, Tim Chown, Lorenzo Colitti, Alejandro D'Egidio, Gert Doering, Fernando Frediani, Guillermo Gont, Steinar Haug, Nick Hilliard, Lee Howard, Christian Huitema, Sheng Jiang, Benjamin Kaduk, Suresh Krishnan, Warren Kumari, Albert Manfredi, Olorunloba Olopade, Jordi Palet Martinez, Pete Resnick, Michael Richardson, Mark Smith, Job Snijders, Sander Steffann, Tarko Tikan, Ole TrÅ,an, Loganaden Velvindron, Å11ric Vyncke, Robert Wilton, Timothy Winters, Christopher Wood, and Chongfeng Xie for providing valuable comments on earlier draft versions of this document.

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.

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 Å½orÅ¾
6connect

Email: jan@6connect.com

Richard Patterson
Sky UK

Email: richard.patterson@sky.uk

Bernie Volz
Individual Contributor
116 Hawkins Pond Road
Center Harbor, NH 03226
United States of America

Email: bevolz@gmail.com