IPv6 Maintenance (6man) Working Group
Internet-Draft
Updates: 4861, 4862 (if approved)
Intended status: Standards Track
Expires: August 22, 2019

F. Gont
SI6 Networks / UTN-FRH
J. Zorz

February 18, 2019

Reaction of Stateless Address Autoconfiguration (SLAAC) to Renumbering Events
draft-gont-6man-slaac-renum-01

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 period of time, thus resulting in connectivity problems. This document analyzes these problem scenarios, and proposes workarounds to improve network robustness. This document updates RFC4861 and RFC4862 to allow for a more robust reaction to network configuration changes.

Table of Contents

1. Introduction

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.

There are a number of scenarios where this problem may arise. For example, the most common deployment scenario for IPv6 in home networks is that in which a CPE router employs DHCPv6 Prefix Delegation (DHCPv6-PD) [RFC8415] to request a prefix from the Internet Service Provider (ISP), and a prefix belonging to the leased prefix is advertised on the LAN-side of the CPE via Stateless Address Autoconfiguration (SLAAC) [RFC4862]. In scenarios where the CPE router crashes and reboots, the CPE may be leased (via DHCPv6-PD) a different prefix from the one previously leased, and will therefore advertise (via SLAAC) the new prefix on the LAN side. Hosts will normally configure an address for the new prefix, but will normally retain and actively employ the previously-configured addresses, since their associated Preferred Lifetime and Valid Lifetime allow them to do so.

Lacking any explicit signaling to "obsolete" the previously-configured addresses (for the now invalid/stale prefix), hosts may continue employing the previously-configured addresses which will typically result in packets being blackholed -- whether because of egress-filtering by the CPE or ISP, or because responses to such packets will be discarded or routed elsewhere.

The default values for the "Valid Lifetime" and "Preferred Lifetime" of Prefix Information Options (PIOs), as specified in [RFC4861], are:

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

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

This means that in the aforementioned scenarios, the stale addresses would be retained for unacceptably long period of time, thus leading to interoperability problems, instead of nodes transitioning to the newly-advertised prefix(es) in a timelier manner.

Some devices have implemented mechanisms to address this problem, such as sending RAs to invalidate the apparently stale prefixes when the device receive any packets employing a source address from a prefix not being advertised for address configuration [FRITZ]. However, this may introduce other interoperability problems, particularly in multihomed scenarios. This is yet another indication that advice in this area is warranted.

This unresponsiveness is caused by the both 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 is desirable that behave in a way that hosts are explicitly notified when configuration information has become stale. However, for robustness reasons it is paramount for hosts to be able

to recover from stale configuration information even if somehow the network is unable to explictly notify hosts about such condition.

Section 3 analysis this problem in more detail. Section 5 proposes workarounds to improve network robustness.

2. Terminology

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

3. Analysis of the Problem

As noted in Section 1, the problem discussed in this document is associated with sub-optimal behaviour and policies of different entities. Each of the following sections analyze each of them in detail.

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

Many protocols, from different layers, normally employ timers. The general logic is as follows:

- A timer is set with a value that, under normal conditions, the timer does *not* go off.

- Whenever a fault condition arises, the timer goes off, and the protocol can perform fault recovery

One common example for the use of timers is when implementing reliability mechanisms where a packet is transmitted, and unless a response is received, a retransmission timer will go off to trigger retransmission of the original packet.

For obvious reasons, the whole point of using timers is that in problematic scenarios, they will go off, and trigger some recovery action.

However, IPv6 SLAAC employs, for PIOs, these default values:

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

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

Under normal network conditions, these timers will be reset/refreshed to the default values. However, 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 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 is 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.

The default values for these timers should be such that, under normal circumstances (including some packet loss), the timers will be refreshed/rest, 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 un-preferring the corresponding addresses and subsequently removing them).

In the context of [RFC8028], where it is clear that use of addresses configured for a given prefix is mostly tied to using the router that advertised the prefix as a next hop, one could tell that neither the "Preferred Lifetime" or the "Valid Lifetime" of a PIO should ever be longer than the default value for the "Router Lifetime" of Router Advertisement packets. This means that since [RFC4861] specifies the default value for the "Router Lifetime" as: AdvDefaultLifetime: 3 * MaxRtrAdvInterval, and MaxRtrAdvInterval defaults to 600 seconds, the values employed for the "Preferred Lifetime" (AdvPreferredLifetime) and "Valid Lifetime" (AdvValidLifetime) of PIOs should never be larger than 1800 seconds (30 minutes). Capping the lifetimes in PIOs as suggested would not eliminate the problem discussed in this document, but certainly reduces the amount of time it takes for hosts to converge to updated network configuration information.

Section 5.1.1 of this document updates the SLAAC specification to employ more appropriate timer values.

3.2. Inability of SLAAC Hosts to Recover from Stale Network Configuration Information

In the absence of explicit signalling from SLAAC routers (such as sending PIOs with a "Valid Lifetime" set to 0), SLAAC hosts fail to recover from stale configuration information in a timely manner. This is exacerbated by the inappropriate timers employed for the lifetime of PIOs, as discussed in Section 3.1.

It is possible to heuristically infer that network configuration information has changed. For example, if the same SLAAC router (as

identified by its link-local address) is advertising new prefixes via PIOs, but has ceased to advertise the previously-advertised prefixes, this should be considered an indication that network configuration information has changed. Implementing 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.

Section 5.1.3 of this document specifies a local policy that SLAAC hosts can implement to recover from stale prefixes.

3.3. 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), such that there is explicit signaling to SLAAC hosts to remove the stale information (including configured addresses and routes).

In order to cope with the possibility of a SLAAC router crashing and rebooting without any state of the previously-advertised prefixes, a SLAAC router should record on stable storage the information of which prefixes were being advertised on which interfaces, such that upon reboots, such prefixes may be advertised with appropriate lifetimes (both "Preferred Lifetime" and "Valid Lifetime" set to 0) to cause hosts to remove any configuration information derived from previous announcements of such prefixes.

Explicit signaling of network configuration changes would eliminate the problem discussed in this document. However, since it is impossible for a host to know whether such explicit signals will be received, they are not relieved from inferring changes in network configuration information, as discussed in Section 3.2.

Section 5.1.2 updates the SLAAC specification such that routers explicitly signal stale configuration to SLAAC hosts. Section 5.3 specifies the corresponding requirements for CPE routers.

3.4. Under-specified 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 specified in two different software components with the interface between the two implemented by some sort of script that takes feeds the SLAAC implementation with values learned from DHCPv6-PD.

Quite frequently, 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. The lifetime values employed for the "Preferred Lifetime" (AdvPreferredLifetime) and "Valid Lifetime" (AdvValidLifetime) should never be larger than the remaining lease time for the corresponding prefix (as learned via DHCPv6-PD).

Section 5.2 of this document specifies this aspect of the interaction between DHCPv6-PD and SLAAC.

4. Use of Dynamic Prefixes

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:

- Provisioning systems may be unable to deliver stable IPv6 prefixes.

- 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.

- Applicable legislation may require the ISP to deliver dynamic IPv6 prefixes *by default* (see e.g. [GERMAN-DP]).

The authors of this document understand that, for a number of reasons (such as the ones stated above), IPv6 deployments may employ dynamic prefixes, or there might be scenarios in which the dynamics of a network are such that the network exhibits the behaviour of dynamic prefixes. Rather than try to preclude how operators may run their networks, this document aims at improving network robustness in the deployed Internet.

5. Possible workarounds

The following subsections discuss possible workarounds for the aforementioned problem. Section 5.1 specifies modifications to SLAAC which include the use of more appropriate lifetime values and a mechanism for hosts to infer when a previously-advertised prefix has become stale. This modification leads to more robust behaviour even for existing deployments.

Section 5.3 specifies the interaction between DHCPv6-PD and SLAAC, such that devices such as CPEs may be in a better position to convey current network information to hosts on the LAN-side. For obvious reasons (legacy CPEs, etc.), this improved behaviour cannot be relied upon for mitigating the problem described in Section 1. However, it does contribute to more robust IPv6 networks.

Finally, Section 4 analyzes the trade-offs of employing stable vs. dynamic network prefixes.

5.1. Improvements to SLAAC

5.1.1. Default Timer Values

5.1.1.1. Router Configuration Variables

The "default" value for the router configuration variable "MaxRtrAdvInterval" (Section 6.2.1 of [RFC4861] is changed to 300 seconds (5 minutes).

As a result of this change, the default values for two other configuration variables are indirectly modified (since they are specified in relation with MaxRtrAdvInterval):

- MinRtrAdvInterval: 0.33 * MaxRtrAdvInterval = 99 seconds

- AdvDefaultLifetime: 3 * MaxRtrAdvInterval = 900 seconds (15 minutes)

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

> AdvValidLifetime: AdvDefaultLifetime (which according to this specification defaults to 900 seconds / 15 minutes)
>
> AdvPreferredLifetime: 0.50 * AdvValidLifetime (which would result in 450 seconds, that is, 7.5 minutes)

The motivation of this update is as follows:

- Link the lifetimes associated with prefixes to the lifetime of the router advertising the prefixes, since use of advertised prefixes is closely tied to the router that has advertised them (as per [RFC8028]).

- Lacking RAs that refresh information, addresses configured for such prefixes would become un-preferred, and thus Rule 3 of [RFC6724] would cause other configured addresses (if available) to be preferred and used instead.

- Limit the amount of time that a stale prefix needs to be advertised with a lifetime od "0" on the local network (see Section 12 of [RFC4861]).

- Employ default router lifetimes that improve the ability of hosts to recover from fault scenarios.

5.1.1.2. Processing of PIO Lifetimes at Hosts

Hosts should cap the "Valid Lifetime" and "Preferred Lifetime" of PIOs to the "Router Lifetime" value in the received Router Advertisement. That is, if the "Valid Lifetime" or "Preferred Lifetime" of a PIO is larger than the "Router Lifetime" value of the Router Advertisement carrying the PIO, the corresponding value should be capped to that of the "Router Lifetime" value of the received RA.

The motivation for this update is as follows:

- Limits the amount of time a host is required to maintain possibly stale information.

5.1.2. Signaling Stale Configuration Information

In order to phase-out stale configuration information, the SLAAC specification is updated as follows:

- A router sending RAs that advertise dynamically-learned prefixes (e.g. via DHCPv6-PD) on an interface MUST record, on stable storage, the list of prefixes being advertised on each network segment.

- Upon changes to the advertised prefixes, and after bootstrapping, the router advertising prefix information via SLAAC should proceed as follows:

  * Any prefixes that were previously advertised via SLAAC, but that are not currently intended for address configuration, MUST

be advertised with a PIO option with the "A" bit set to 1 and the "Valid Lifetime" and a "Preferred Lifetime" set to 0.

* Any prefixes that were previously advertised via SLAAC as "on-link", but that are not currently not considered "on-link", MUST be advertised with a PIO option with the "L" bit set to 1 and the "Valid Lifetime" and a "Preferred Lifetime" set to 0.

* If both of the previous conditions are met (a prefix was previously advertised with both the "A" and "L" bits set, but is currently *not* intended for address configuration and is *not* considered on-link), the prefix MUST be advertised with a PIO option with both the "A" and "L" bits set to 1 and the "Valid Lifetime" and a "Preferred Lifetime" set to 0. That is. the advertisements of the previous two steps can be coalesced into a single one with both the "A" and "L" bits set.

* The aforementioned advertisement SHOULD be performed for at least the "Valid Lifetime" previously employed for such prefix.

Additionally, 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.

to:

> 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.

The rationale for this update is:

o 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.

5.1.3. Recovery from Stale Configuration Information

The goal of the mechanism 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.

The basic premise behind the mechanism specified in this section 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, configured addresses for the stale prefixes are initially "un-preferred" (such that they are not employed for new communication instances), and they are subsequently removed (if this condition persists).

Local information maintained for each prefix advertised by each router is augmented with one boolean flag named "LTA" (Lifetime Avoidance) that defaults to "FALSE". 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].

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

- The LTA flag for each prefix advertised in the current Router Advertisement should be set to "FALSE".

- For each prefix that had been previously advertised by this router but that does not have a corresponding PIO with the "A" flag set in the received RA, proceed as follows:

  * IF the LTA flag is "TRUE", then:

    + IF there is any address configured for this prefix with a "Preferred Lifetime" larger than 0, set its "Preferred Lifetime" to 0, and the LTA flag of this prefix to "FALSE".

    + ELSE (all addresses for this prefix have a "Preferred Lifetime" of 0), set the "Valid Lifetime" of all addresses configured for this prefix to 0, and the LTA flag of this prefix to "FALSE". This will cause removal of all addresses for this prefix. Additionally, if the corresponding prefix had been advertised as on-link ("L"=1) by this router, remove any routes to this prefix associated with the network interface card on which the RA packet was received.

  * ELSE (the LTA flag is "FALSE"):

    + Set the LTA flag of the corresponding prefix to "TRUE".

Appendix B illustrates the packet exchange and operation of the algorithm for a typical scenario. Appendix A provides a flowchart for this algorithm.

NOTES:

- The aforementioned processing assumes that while network configuration information might be split into multiple RAs, PIOs will be spread among *at most* two RAs. This assumption avoids the use of any timers for this specific purpose.

- 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 new prefix being advertised, the traditional processing is unaffected (the mechanism discussed in this document will *never* be triggered because no packets with PIOs with the "A" flag will be received). The logic here is that it is better to have some address, than no address at all.

- The processing of RAs that do not contain any PIOs with the "A" bit set remains unaffected.

- The specified modification takes the conservative approach of first setting the "Preferred Lifetime" to 0 (such that addresses become non-preferred), and subsequently setting the "Valid Lifetime" to 0 (such as the addresses are completely removed). Once the addresses for this prefix have been removed, routes to this prefix associated with the network interface on which the RA packets were received are also removed.

- 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", and less than five minutes to completely remove such addresses from the system. 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, 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 * MaxRtrAdvInterval) and MaxRtrAdvInterval (default: 600 seconds) has elapsed, thus meaning that the stale addresses would be removed between 3.3 and 10 minutes of being marked as "not preferred".

5.2. Interaction Between DHCPv6-PD and SLAAC

The "Preferred Lifetime" and "Valid Lifetime" of PIOs corresponding to prefixes learned via DHCPv6-PD MUST NOT span past the lease time of the DHCPv6-PD prefixes. This means that the advertised "Preferred Lifetime" and "Valid Lifetime" MUST be dynamically adjusted such that the advertised lifetimes never span past the lease time of the prefixes delegated via DHCPv6-PD.

This is in line with these existing requirements from other specifications, which we reference here for clarity:

- [RFC8415] specifies, in Section 6.3, that "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."

5.3. Improved CPE behavior

This section specifies and clarifies requirements for CPE routers (particularly when they advertise prefixes learned via DHCPv6-PD) that can help mitigate the problem discussed in Section 1. This improves the situation for hosts that do not implement the modification specified in Section 5.1 but would obviously make robustness dependent on the CPE (on which the user or ISP may have no control), as opposed to the host itself. This mechanism is mostly orthogonal to the improved host behaviour specified in Section 5.1.

The updated behaviour is as follows:

- The CPE MUST signal stale configuration information as specified in Section 5.1.2

- The CPE MUST implement the DHCPv6-PD/SLAAC interface specified in Section 5.2.

The aforementioned improved behaviour assumes compliance with the following existing requirements from other specifications, which we reference here for clarity:

- [RFC7084] specifies (requirement LE-13, in Section 4.3) that when the delegated prefix changes (i.e., the current prefix is replaced with a new prefix without any overlapping time period), "the IPv6 CE router MUST immediately advertise the old prefix with a Preferred Lifetime of zero and a Valid Lifetime of either a) zero or b) the lower of the current Valid Lifetime and two hours (which

must be decremented in real time) in a Router Advertisement message as described in Section 5.5.3, (e) of [RFC4862]"

6. IANA Considerations

This document has no actions for IANA.

7. Security Considerations

This document discusses a a problem that may arise in scenarios where dynamic IPv6 prefixes are employed, and proposes workarounds that enable such usage while avoiding interoperability problems. The security and privacy implications of IPv6 addresses are discussed in [RFC7721].

An attacker that could impersonate a router could forge multiple RA packets that contain PIOs of prefixes that are currently not advertised on the local network, to trigger the mechanism specified in this document to cause addresses currently configured for the legitimate prefixes to be removed. However, an attacker that can impersonate a router could more easily achieve the same goal by advertising the legitimate prefixes with both the "Preferred Lifetime" and "Valid Lifetime" set to 0.

Capping the "Valid Lifetime" and "Preferred Lifetime" at hosts limit the length of the effects of a non-sustained attack, since hosts would now be able to recover in a timelier manner.

Attacks based on forged RA packed can be mitigated with technologies such as RA-Guard [RFC6105] [RFC7113].

8. Acknowledgments

The authors would lie to thank (in alphabetical order) Mikael Abrahamsson, Luis Balbinot, Brian Carpenter, Gert Doering, Nick Hilliard, Philip Homburg, Lee Howard, Christian Huitema, Albert Manfredi, Jordi Palet Martinez, Richard Patterson, 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.

The problem discussed in this document has been previously documented in [RIPE-690].

9. References

9.1. 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>.

[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>.

[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>.

9.2. Informative References

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

[GERMAN-DP] BFDI, "Einfuhrung von IPv6 Hinweise fur Provider im Privatkundengeschaft und Herstellere", Entschliessung der 84. Konferenz der Datenschutzbeauftragten des Bundes und der Lander am 7./8. November 2012 in Frankfurt (Oder), November 2012, <http://www.bfdi.bund.de/SharedDocs/Publikationen/Entschliessungssammlung/DSBundLaender/84DSK_EinfuehrungIPv6.pdf?__blob=publicationFile>.

[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>.

[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>.

[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>.

[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>.

[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>.

Appendix A. Flowchart for Host Processing of RAs

Conceptually, the mechanism operates as follows:

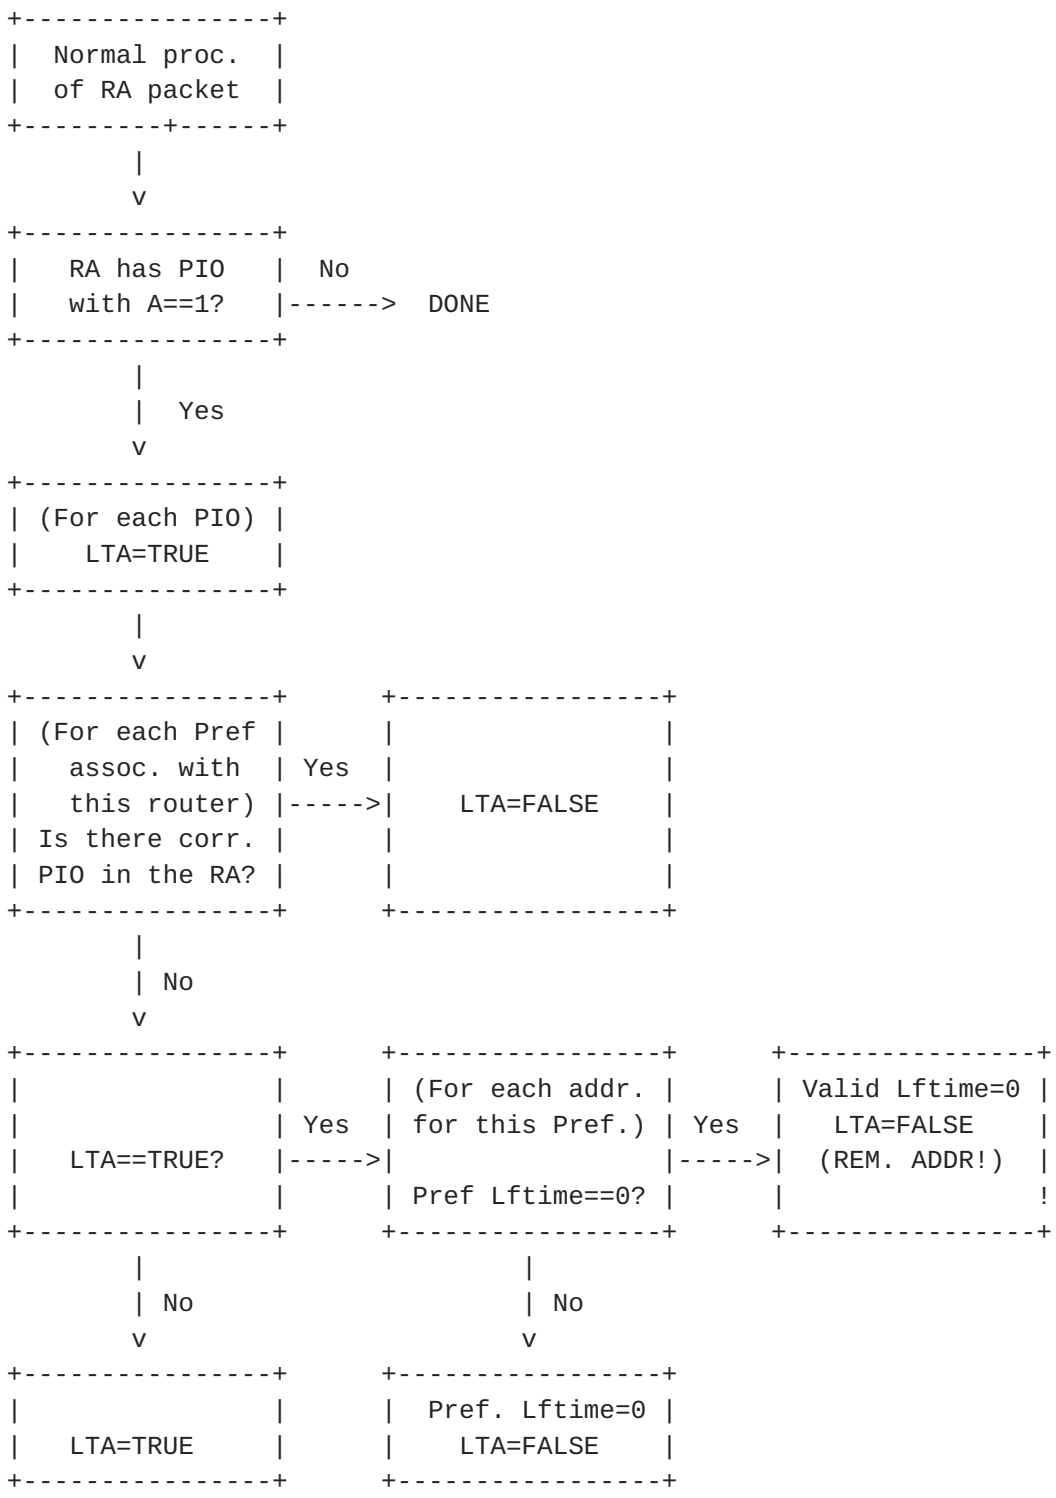

Appendix B. Sample Timeline for Host Processing of RAs

The following example illustrates a sample packet exchange that illustrates the algorithm specified in Section 5.1:

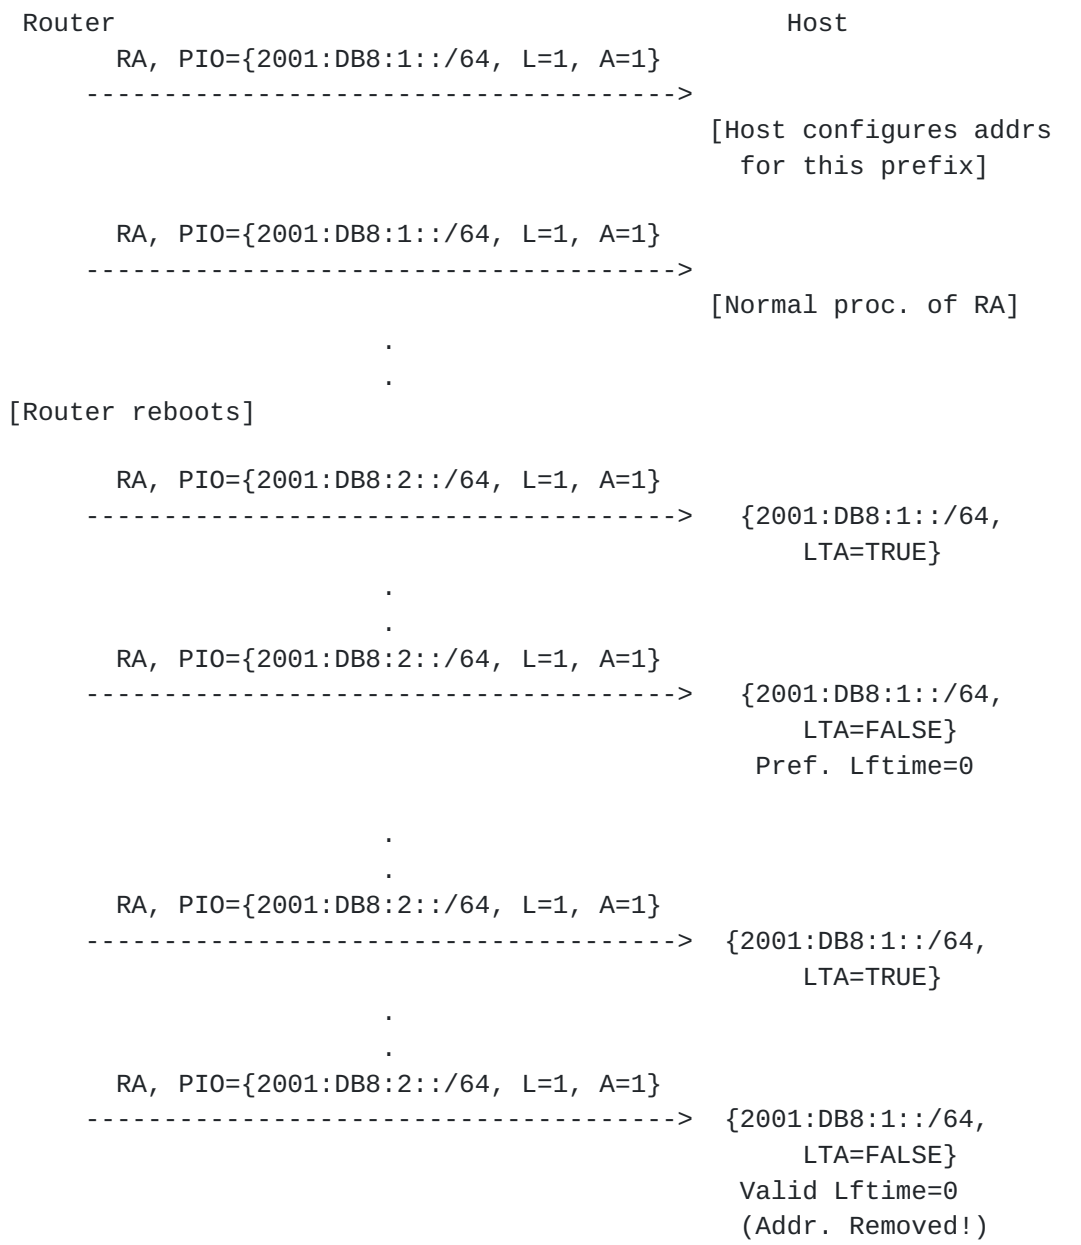

Appendix C. Analysis of Some Suggested Workarounds

[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 (e.g. on the 6man list). The following subsections analyze these suggested workarounds, in the hopes of avoiding rehashing discussions of such workarounds.

C.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]) might 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 might perform heuristics that might help to mitigate the problem discussed in this document.

The aforementioned proposal has a number of drawbacks and limitations:

- It assumes that the CPE enforces ingress/egress filtering [RFC2827]. While this is desirable behaviour, it cannot be relied upon.

- 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.

- ICMPv6 Type 1 Code 5 could be interpreted as the employed address to be 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.

- Reacting to these error messages would create a new attack vector. This is of particular concern since ICMP-based attack do not even require that the Source Address of the attack packets be spoofed [RFC5927].

C.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.

It has been suggested that 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:

- 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.

- 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.

- [RFC6724] is updated such that this rule is incorporated:

  Rule 3.5: 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).

There are a number of limitations and drawbacks associated with this approach:

- It may help to improve the selection of a source address, but it does not help with the housekeeping of configured information:

  * If the stale prefix is re-used in another network, nodes employing stale addresses and routes for this prefixes will be unable to communicate with the new "owners" of the prefix.

  * Given that 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.

- The earlier the above rule is incorporated into the [RFC6724] rules, the more it could lead to suboptimal address selection. For example, if incorporated as Rule 3.5 (before Rule #4, but after Rule 3), an implementation may end up choosing a source address from a "fresher" prefix rather than one with a longest-matching prefix, thus leading to suboptimal routing. On the other hand, the later the rule is incorporated into the [RFC6724] rules, the more irrelevant the rule becomes (since existing rules might prefer the stale addresses).

- In scenarios where multiple prefixes are being advertised (whether by a single router via multiple RAs, multiple routers on the same LAN segment, or different routers on different LAN segments (when a host has multiple interfaces)), the new SAS rule is guaranteed to result in non-deterministic behaviour, with hosts frequently

changing the selected address. This is certainly not desirable from a troubleshooting perspective.

As a result, updating IPv6 source address selection does not relieve nodes from improving their SLAAC implementations as specified in Section 5.1, if at all desirable. On the other hand, employing appropriate timers as specified in Section 5.1.1 would result in Rule 3 of [RFC6724] employing fresh addresses, without leading to undeterministic behaviour.

Authors' Addresses

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

Phone: +54 11 4650 8472
Email: fgont@si6networks.com
URI: https://www.si6networks.com

Jan Zorz

Email: jan@go6.si