Workgroup: IPv6 operations
Internet-Draft: draft-link-v6ops-gulla
Published: 25 February 2024
Intended Status: Informational
Expires: 28 August 2024
Authors: J. Linkova
Google

Using Subnet-Specific Link-Local Addresses to Improve SLAAC Robustness

Abstract

This document suggests that a link-local address used by a router as a source address for Router Advertisement packets is calculated as a function of prefixes listed in the Prefix Information Option of the Router Advertisement. The proposed approach, combined with the Rule 5.5 of the Default Source Address Selection algorithm (RFC6724) and first-hop selection requirements for hosts (RFC 8028) improves the robustness of the SLAAC by allowing the hosts to detect the IPv6 subnet changes much faster and select the correct source address.

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 28 August 2024.

Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction

IPv6 Stateless Address AutoConfugration (SLAAC, [RFC4862] provides IPv6 hosts with a mechanism to configure their IPv6 stack based on the information (such as an IPv6 prefix and the default router address) provided by the on-link routers. If that information changes (e.g. a prefix assigned to the link is changed), the routers need to explicitly invalidate the outdated information (e.g. by sending a Router Advertisement packet which deprecates the old prefix). In the absence of an explicit signal the host would be using the outdated information until its lifetime expires. Multiple documents discuss the SLAAC renumbering problem and proposed various improvements to the host and router behaviour (see [RFC9096] and draft-ietf-6man-slaac-renum).

This document recommends that the link-local address the router sends the router advertisement from should depend on the network prefix(es) assigned to the router interface. As a result, Router Advertisements containing different sets of PIOs are sent from different link-local addresses. That allows the hosts to select the source address from the prefix advertized by the reachable next-hop and recover from a renumbering or network segment change events much faster.

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Terminology

PIO: Prefix Information Option, [RFC4861].

RA: Router Advertisement, [RFC4861].

SLAAC: StateLess Address AutoConfiguration, [RFC4862].

4. Subnet Change Scenarios

There are multiple scenarios when an IPv6 subnet assigned to the link can change without any explicit signals received by hosts. When it happens, the hosts can end up with outdated IPv6 address configuration, which leads to broken connectivity and degraded user experience. Examples include but are not limited to:

* A prefix delegated to a CPE router changes, and the router does not send RAs to the connected hosts, deprecating the old prefix. The hosts would be using addresses from both old and new prefixes until the prefix lifetime expires.

* A host is connected to a wired port, and a VLAN (and the corresponding IPv6 subnet) changed. This often happens if the VLAN is configured as a result of 802.1x or MAC-based authentication, so the VLAN is provided by RADIUS. Some OSes do not correctly reset IPv6 stack when the wired interface 802.1x state changes from ‘unauthenticated’ to ‘authenticated’ (or vice versa).

In all those scenarios a host might move between IPv6 subnets without complete disconnection and without detecting the network change. As a result the following sequence of events may occur, leading to broken connectivity:

* The host is connected to a network A, receives an RA from the router with a PIO containing pref_a, forms IPv6 addresses from that prefix using SLAAC.

* The host attachment changes from network A to network B or an IPv6 prefix configured on the network changes from pref_a to pref_b. The host doesn’t detect the network change and doesn’t clear the IPv6 stack.

*The host receives an RA from the router with a new PIO for pref_b and forms new addresses from that prefix.

*Now the host has two sets of IPv6 addresses - one from pref_a and one from pref_b. Addresses from pref_a are unusable: even if the outgoing packets are not dropped by anti-spoofing filters,the return traffic wouldn’t be able to reach the host. So if the host selects an address from pref_a as a source address for outgoing communication (as per RFC6724 or by using any other custom algorithms), the traffic would be dropped, causing user-visible outages.

It should be noted that the Detecting Network Attachment algorithm defined in [RFC6059] relies on the combination of the link-layer address and the link-local IPv6 address of a router to be unique across links. However that assumption doesn't often hold. For example, network administrators prefer to configure the same, easy to remember link-local address (e.g. 'fe80::1') to router interfaces on different links. Some router implementations are known to use the virtual router MAC address to create the Modified Extended Unique Identifier (EUI)-64 identifiers for VRRPV3 virtual link-local addresses (violating Section 7.4 of [RFC5798]). As a result, all links with the same VRRP ID (and hence the same virtual router MAC address) would have the same virtual link-local address as well.

5. Default Address Selection Rule 5.5, Default Router Selection and Renumbering

Rule 5.5 of the Default Source Address Selection ([RFC6724]) requires the host to prefer addresses in a prefix advertised by the next-hop. It allows the multihomed host to select the source address correctly: when two routers advertize different prefixes, the host will be sending packets with source address from a given prefix to the router the prefix was received from.

In case of renumbering if both old and new prefixes are advertized by the same router (received from a router with the same link-local address), then Rule 5.5 doesn't help selecting the correct (working) source address. However if the subnet change also leads to the default router address change, then a host implementing Rule 5.5 could recover from the renumbering quickly:

*The host is connected to a network A, receives an RA from the router (link-local address LLA_A) with a PIO containing pref_a, forms IPv6 addresses from that prefix using SLAAC.

*The host attachment changes from network A to network B or an IPv6 prefix configured on the network changes from pref_a to

pref_b. The host doesn’t detect the network change and doesn’t clear the IPv6 stack.

*The host receives an RA from the router (link-local address LLA_B) with a new PIO for pref_b and forms new addresses from that prefix.

*Link-local address LLA_A is not reachable anymore, as the host changes the network attachment point. Neighbor Unreachability Detection ([RFC4861]) detects it and removes LLA_A from the list of default routers.

*The host is using LLA_B as a next-hop for outgoing traffic, so addresses from the pref_b are selected, and addresses from pref_a are not used.

As per [RFC8028], "A host SHOULD select default routers for each prefix it is assigned an address in. Routers that have advertised the prefix in their Router Advertisement message SHOULD be preferred over routers that do not advertise the prefix.". If the host complies with [RFC8028], then the proposed mechanism would work even better, and would provide fast recovery from a renumbering event.

*The host selects a default router (LLA_A) for pref_a.

*When the host receives an RA from LLA_B, containing pref_b, the host selects another default gateway, LLA_B.

*Neighbor Unreachability Detection detects that LLA_A is not reachable, and removes it from the neighbor cache table, so the host can not use it as a default gateway anymore. The host switches to LLA_B and, in accordance with Rule 5.5, starts using addresses from pref_b.

6. Outage Duration During Renumbering Event

When the IPv6 subnet changes (either because the given link has been renumbered, or because the client has moved to another link), there are two factors contributing to the duration of the outage:

*Time required for the host to receive new configuration information (RAs containing new PIOs).

*Time required for the host to deprecate the old configuration information.

6.1. Receiving New Configuration Information from Routers

There is always a delay between the network subnet change (renumebring event or link change event) and arrival of an RA. To

reduce that delay in case of undetected link change, the network administrators SHOULD reduce MinRtrAdvInterval and MaxRtrAdvInterval ([RFC4861]) to ensure that RAs are sent often. In case of the link renumbering, the router SHOULD notify hosts by sending an RA with the new information. While Requirement L-13 of [RFC7084] requires that the IPv6 CE router MUST send an RA immidiately after the delegated prefix changes, that section does not explicitly requires the new prefix to be included. Also, that behaviour is only mandatory for CPE devices, while other routers (Enterprise-grade devices, for example) comply with Section 6.2.4 of [RFC4861] which says:

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

To reduce the time required for hosts to detect renumbering event, it might be desirable to upgrade that "MAY" to "SHOULD".

6.2. Hosts Deprecating the Outdated Configuration

Without changes proposed in this document, a host might be using the outdated prefix for the duration of the PIO preferred lifetime. As per [RFC4861], the default value for preferred lifetime is 604800 secs (7 days). The expired draft draft-ietf-6man-slaac-renum proposes to reduce that value to 14400 seconds (4 hours).

The solution proposed in this document allows hosts which implement Rule 5.5 of the source address selection ([RFC6724]) to stop using the outdated prefix much faster. The time required for the host to detect that the old prefix shouldn't be used for initiating new session is the time required for Neighbor Unreachability Detection (NUD, [RFC4861]) to remove an unreachable entry for the old link-local address of the default router. The default value would be (without taking randomisation factors into account): ReachableTime milliseconds (to move from REACHABLE to STALE) + DELAY_FIRST_PROBE_TIME + MAX_UNICAST_SOLICIT*RetransTimer = 30 seconds + 5 second + 3*1 = 38 seconds.

7. Generating Subnet-Specific Link-Local Addresses for Router Interfaces

*The router SHOULD send router advertisement packets from a dedicated link-local address.

* That dedicated link-local address SHOULD change if the set of prefixes advertized in the Router Advertisement changes. In other words, when the set of prefixes advertized on a given interface changes, the router SHOULD generate a new link-local address and use that address as a source address for Router Advertisements. As soon as a new link-local address is generated, the router SHOULD transmit Router Advertisements as specified in Section 6.2.4 of [RFC4861].

* The router SHOULD stop responding to Neighbor Solicitation packets to the old link-local address. This is required for Neighbor Unreachability Detection mechanism on hosts to mark the old address as unreachable and make the hosts to use the new address as a default router.

Routers which act as DHCPv6-PD clients need to implement some algorithm to generate the interface ID based on the set of prefixes advertised on a given interface. The exact algorithm is outside of scope of this document - it could be some form of hashing function which consumes a list of network prefixes and generates a 64-bits interface ID. For example, the router MAY use the algorithm defined in [RFC7217] and use the list of PIOs as Network_ID.

If the interface subnets are configured statically, the network administrator can configure link-local addresses statically as well. In some cases it might be possible to just utilize the global subnet prefix as an interface ID. For example, if the router has two interfaces configured with 2001:db8:1:1::/64 and 2001:db8:2:2::/64 subnets respectively, it is sufficient to configure fe80::2001:db8:1:1 and fe80::2001:db8:2:2 as the link-local addresses for those interfaces. If the first-hop redundancy is provided by VRRPv3, there is no need to configure the static link-local interface addresses but the virtual link-local address SHOULD be configured instead.

Discussion: if a given router's interface has multiple prefixes configured, it's possible than only one prefix changes. For example:

* An interface has both GUA and ULA prefixes configured, and the global prefix delegated via DHCPv6-PD changes, while ULA stays the same.

* In a multihomed environment, an interface migth have two prefixes delegated by two upstream networks. Those prefixes can change independently.

If all PIOs are advertised in a single RA, and the link-local address used as an RA source, is generated as described in this document, changing just one prefix would impact traffic sent from

all addresses (as the previous default gateway becomes unreachable). Therefore it might be desirable for a router to send multiple RAs - one per PIO, and use a dedicated PIO-specific source address for each of those RAs. In that case, if the host complies with [RFC8028] and selects default routers for each prefix it is assigned an address in, then only traffic from source addresses in the renumbered prefix is impacted (as only the default gateway for that prefix becomes unreachable and needs to be reselected). On the other hand such proposal would lead to increased numbers of RAs being sent, which might negatively impact hosts battery life. It should be noted that an expired draft draft-ietf-6man-slaac-renum recommends against such behaviour and requires all PIOs to be advertized in a single RA.

7.1. Subnet-Specific and Stable Link-Local Addresses

In many cases it might be beneficial for a router to have a stable link-local address (e.g. if that address is advertized as a DNS server, or for management purposes. Router MAY generate subnet-specific link-local addresses in addition to a stable link-local address. Alternatively, if the router is not capable of supporting multiple link-local addresses, it MAY generate a new stable link-local address every time the set of prefixes on the interface changes.

It should be noted that the proposed mechanism assumes that the router does not use the modified EUI-64 format for generating interface ID. As per Section 3 of [RFC8064], nodes SHOULD NOT use the modified EUI-64 format, and SHOULD use the algorithm defined in [RFC7217] instead.

8. Security Considerations

To be added.

9. Privacy Considerations

This document does not introduce any privacy considerations.

10. IANA Considerations

This memo does not introduce any requests to IANA.

11. References

11.1. 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>.

[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>.

[RFC6877] Mawatari, M., Kawashima, M., and C. Byrne, "464XLAT: Combination of Stateful and Stateless Translation", RFC 6877, DOI 10.17487/RFC6877, April 2013, <https://www.rfc-editor.org/info/rfc6877>.

[RFC7050] Savolainen, T., Korhonen, J., and D. Wing, "Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis", RFC 7050, DOI 10.17487/RFC7050, November 2013, <https://www.rfc-editor.org/info/rfc7050>.

[RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <https://www.rfc-editor.org/info/rfc7217>.

[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>.

[RFC8781] Colitti, L. and J. Linkova, "Discovering PREF64 in Router Advertisements", RFC 8781, DOI 10.17487/RFC8781, April 2020, <https://www.rfc-editor.org/info/rfc8781>.

[RFC8925] Colitti, L., Linkova, J., Richardson, M., and T. Mrugalski, "IPv6-Only Preferred Option for DHCPv4", RFC 8925, DOI 10.17487/RFC8925, October 2020, <https://www.rfc-editor.org/info/rfc8925>.

11.2. Informative References

[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>.

[RFC5798] Nadas, S., Ed., "Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6", RFC 5798, DOI 10.17487/RFC5798, March 2010, <https://www.rfc-editor.org/info/rfc5798>.

[RFC6059] Krishnan, S. and G. Daley, "Simple Procedures for Detecting Network Attachment in IPv6", RFC 6059, DOI 10.17487/RFC6059, November 2010, <https://www.rfc-editor.org/info/rfc6059>.

[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>.

[RFC8064] Gont, F., Cooper, A., Thaler, D., and W. Liu, "Recommendation on Stable IPv6 Interface Identifiers", RFC 8064, DOI 10.17487/RFC8064, February 2017, <https://www.rfc-editor.org/info/rfc8064>.

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

Acknowledgements

Thanks to Dale W. Carder, Brian Carpenter, Lorenzo Colitti, Fernando Gont, Alexandre Petrescu, Mark Smith, Ole Troan, Eduard Vasilenko, Eric Vyncke for the discussions, the input and all contribution.

Author's Address

Jen Linkova
Google
1 Darling Island Rd
Pyrmont NSW 2009
Australia

Email: furry13@gmail.com, furry@google.com