IPv6 Operations Working Group (v6ops)
Internet-Draft
Intended status: Informational
Expires: 4 September 2025

F. Gont
SI6 Networks

3 March 2025

Problem Statement about IPv6 Support for Multiple Routers, Multiple Interfaces, and Multiple Prefixes
draft-gont-v6ops-multi-ipv6-02

Abstract

This document discusses current limitations in IPv6 Stateless Address Auto-configuration (SLAAC) that prevent support for common multi-router, multi-interface, and multi-prefix scenarios. It provides discussion on the challenges that these scenarios represent, and why a solution in this space is warranted. Finally, it specifies a number of common scenarios that any solution in this space should be able to 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 4 September 2025.

Copyright Notice

Table of Contents

1. Introduction

IPv6 Stateless Address Autoconfiguration (SLAAC) [RFC4862] is based on the assumption that SLAAC routers advertise configuration information on a local network, and SLAAC hosts will aggregate this information and use it as they see fits. Im simple network scenarios where there is a single local router, or where there are multiple routers but all such routers advertise the same network configuration information and provide the same service, SLAAC works just fine. However, other more complex (yet very common) scenarios are currently unsupported. These scenarios include:

* A host attaches to a local-area network (LAN) that employs two different routers, one for each upstream Internet Service Provider (ISP).

* A host attaches to two (or more) different networks via two (or more) network interfaces.

* A host attaches to a local-area network, and receives conflicting information from two or more routers.

In the first two scenarios, a SLAAC host ends up receiving information from two routers that are managed by different entities (two different ISPs) and, for all practical purposes, each piece of configuration information advertised by each router is only of use when employed in conjunction with the rest of the information advertised by such router. In other words, mixing configuration information from different SLAAC advertising routers will usually lead to interoperability problems.

The third scenario could be considered a corner-case of the first two scenarios: two or more routers send conflicting information, such as the same SLAAC configuration information with different lifetimes (e.g., one SLAAC router advertises a piece of information with a lifetime of zero, and another advertises the same information with non-zero lifetime). In this scenario, a single router advertising configuration information with a lifetime of zero may simply cause the corresponding information to be e.g. completely discarded from the host.

Section 2 defines the terminology employed throughout this document. Section 3 elaborates on a number of scenarios that are generally not supported by current implementations, which not only serve to illustrate the problem statement, but also as test cases that any solution in this space should be able to address gracefully. Section 5 discusses future work may be needed to address the problem at hand.

2. Terminology

Multi-prefix scenario:
: A network scenario where a host employs two or more IPv6 prefixes for address configuration with SLAAC. This may be the result of attaching to two or more networks via two or more network interfaces, or simply the result of attaching to a single local network where multiple prefixes are advertised via one or more SLAAC routers.

Multi-router scenario:
: A network scenario where a SLAAC hosts employs two or more SLAAC routers. This may be the result of attaching to two o more networks via two or more network interfaces, or simply the result of attaching to a single local network where multiple prefixes are advertised via one or more SLAAC routers.

Multi-interface scenario:
A network scenario where a host employs two or more network interfaces (without considering the "loopback" interface). Some bibliography refer to these hosts as being "multihomed". In the vast majority of cases, hosts employing multiple interfaces will result in "multi-prefix" and "multi-router" scenarios. In the specific corner case where a host attaches to the same network via two (or more) network interfaces, this will typically result in a multi-router scenario, where the same router that is available via multiple is interfaces is considered, fo all practical purposes, as a different router -- .e.g., the same router will result in different default routes, one via each of the network interfaces.

Multi-IPv6:
The term "Multi-IPv6" (case insensitive) is employed throughout this document as a short-hand for network scenarios where multiple IPv6 routers, multiple interfaces, and/or multiple IPv6 prefixes are employed. We note that "multi-prefix", "multi-interface", and "multi-router" scenarios are not mutually-exclusive: for instance, a host that employs multiple interfaces will usually also employ multiple prefixes and multiple routers.

SLAAC prefix set:
The SLAAC prefix set for an SLAAC router is composed of all prefixes that are being advertised by the router via SLAAC PIOs (irrespective of how e.g. the "A" and "L" flags of the PIO were set).

3. Scenarios

3.1. Multi-Router Scenario

Consider a network scenario where a user attaches two Customer Premises Equipment (CPE) routers to a local network ("Network_C" in our example) for improved network resilience, The scenario could be described as follows:

* Two SLAAC routers (ROUTER_A and ROUTER_B) from different ISPs (ISP_A and ISP_B, respectively) are attached to Network C

* ROUTER_A advertises:

  - Prefix PREFIX_A for address configuration (by means of a Prefix Information Option (PIO) [RFC4861]).

  - One Recursive DNS server (RDNSS_A) (by means of Recursive DNS Server option [RFC8106]).

* ISP_A enforces ingress filtering [RFC2827], and implements ACLs such that RDNSS_A only processes requests from ISP_A customers.

* ROUTER_B advertises:

  - Prefix PREFIX_B for address configuration (by means of a Prefix Information Option (PIO) [RFC4861]).

  - One Recursive DNS server (RDNSS_B) (by means of Recursive DNS Server option [RFC8106]).

* ISP_B enforces ingress filtering [RFC2827], and implements ACLs such that RDNSS_B only processes requests from ISP_B customers.

* Host C attaches to Network C, and thus configures:

  - Addresses from both prefixes (PREFIX_A and PREFIX_B).

  - Two default routers (ROUTER_A and ROUTER_B).

  - Two recursive DNS servers: RDNSS_A and RDNSS_B.

In this scenario, Host C may only send traffic from PREFIX_A via ROUTER_A or from PREFIX_B via ROUTER_B: otherwise, packets will be dropped as a result of ingress filtering [RFC2827]. Similarly, Host C may only send DNS queries from PREFIX_A to RDNSS_A, or from PREFIX_B to RDNSS_B: sending traffic from PREFIX_A to RDNSS_B or from PREFIX_B to RDNSS_A will result in the ACLs enforced by the respective ISPs to drop the DNS queries.

It should be noted that it is quite common for DNS responses to depend on the source address of the query, for improved service. For example, if the "www.example.com" web site was served via a Content Delivery Network (CDN), RDNSS_A would likely resolve that domain name to an IP address that is topologically close to ISP_A, while RDNSS_B would resolve the same domain name to an address that is topologically close to ISP_B. In many cases, the corresponding web cache might be hosted within the ISP network itself. In these scenarios, using the information learned from RDNSS_A via ISP_B, or using the information learned via RDNSS_B via ISP_A would likely lead to sub-optimal service (e.g., larger Round-Trip Time) or no service (if ACLs were being enforced).

It should be evident that each piece of information being advertised via SLAAC is only usable when employed in conjunction with the rest of the information advertised by the same router. However, SLAAC does not require this behavior: hosts are free to use any piece of configuration they learn via SLAAC as they see fit.

As a result of these considerations, in a scenario where e.g. ROUTER_A becomes unreachable while ROUTER_B is known to be reachable:

* RDNSS_A should not be employed for DNS resolution, since queries to RDNSS_A would need to be performed from PREFIX_A via ROUTER_A (considered unreachable). Thus, in this scenario DNS queries should be performed from PREFIX_B via ROUTER_B to RDNSS_B.

* Cached DNS information learned via RDNSS_A should be removed/ ignored, since they would require that new communication instances to the corresponding IPv6 addresses would employ addresses from PREFIX_A via ROUTER_A (considered unreachable).

* New communication instances to an IPv6 literal (i.e., an IPv6 address that has not been learned via the DNS) should not employ ROUTER_A as the next hop (since it is considered unreachable). Hence, in this scenarios new communication instances should employ addresses from PREFIX_B via ROUTER_B.

Section Section 3.2 and Section 3.3 discuss the expected behavior in two specific failure cases.

3.2. Multi-Router/Multi-Prefix Failover Scenario

Consider a network scenario where a user attaches two Customer Premises Equipment (CPE) routers from different ISPs to a local network ("Network_C" in our example) for improved network resilience, The scenario could be described as follows:

* Two SLAAC routers (ROUTER_A and ROUTER_B) from different ISPs (ISP_A and ISP_B, respectively) are attached to Network C

* ROUTER_A advertises:

  - Prefix PREFIX_A for address configuration (by means of a Prefix Information Option (PIO) [RFC4861]).

  - One Recursive DNS server (RDNSS_A) (by means of Recursive DNS Server option [RFC8106]).

* ISP_A enforces ingress filtering [RFC2827], and implements ACLs such that RDNSS_A only processes requests from ISP_A customers.

* ROUTER_B advertises:

  - Prefix PREFIX_B for address configuration (by means of a Prefix Information Option (PIO) [RFC4861]).

- One Recursive DNS server (RDNSS_A) (by means of Recursive DNS Server option [RFC8106]).

* ISP_B enforces ingress filtering [RFC2827], and implements ACLs such that RDNSS_B only processes requests from ISP_B customers.

* Host C attaches to Network C, and thus configures:

  - Addresses from both prefixes (PREFIX_A and PREFIX_B).

  - Two default routers (ROUTER_A and ROUTER_B).

  - One recursive DNS server (RDNSS_A).

Consider the case where e.g. ROUTER_A becomes unreachable. If ROUTER_B is still considered reachable:

* Ongoing communication instances employing source addresses from PREFIX_A via ROUTER_A should continue employing ROUTER_A as the next-hop, since ROUTER_A is the only router advertising PREFIX_A.

* New communication instances should employ addresses from PREFIX_B via ROUTER_B (as next-hop), since ROUTER_B is the only router known to be reachable, and PREFIX_B is the only prefix being advertised by ROUTER_B.

Hosts implementing [RFC8028] will result in ongoing communications from PREFIX_A via ROUTER_A will continue employing ROUTER_A as the next-hop. However, [RFC6724] does not take into consideration the reachability of the next-hop when performing source address determination, and hence may end up selecting a source address that would end up employing an unreachable next hop (an address from PREFIX_A via ROUTER_A as the next hop, in our scenario).

NOTE:
In this scenario, both ROUTER_A and ROUTER_B advertise the same RNDSS (RDNSS_A), and therefore the information learned from that DNS server is presumably equally usable from PREFIX_A via ROUTER_A and from PREFIX_B via ROUTER_B. Hence, if ROUTER_A becomes unreachable while ROUTER_B is still known to be reachable, hosts SHOULD prefer using PREFIX_B via ROUTER_B.

In a scenario where ROUTER_A advertised RDNSS_A, and ROUTER_B advertised RDNSS_B, and ROUTER_A became unreachble, one could envision the following cases:

1. Since ROUTER_A is unreachable, RDNSS_A would be unreachable. Hence RDNSS_B would be employed to resolve domain names, and as per Section 3.1, connections to IP addresses obtained via RDNSS_B would employ addresses from PREFIX_B via ROUTER_B.

2. As per Section 3.1, a host should not employ cached DNS entries obtained via a RDNSS where none the router that advertised RDNSS_B are known to reachable, since that would force the host to employ an unreachable router as the next hop.

3.3. Multi-Router Failover Scenario

Consider the case where two routers attach to the same network (Network_C), and advertise the same configuration information. That is,

* Two SLAAC routers (ROUTER_A and ROUTER_B) are attached to Network_C.

* ROUTER_A advertises:

  - Prefix PREFIX_A for address configuration (by means of a Prefix Information Option (PIO) [RFC4861]).

  - One Recursive DNS server (RDNSS_A) (by means of Recursive DNS Server option [RFC8106]).

* ROUTER_B advertises:

  - Prefix PREFIX_A for address configuration (by means of a Prefix Information Option (PIO) [RFC4861]).

  - One Recursive DNS server (RDNSS_A) (by means of Recursive DNS Server option [RFC8106]).

* Host C attaches to Network_C, and thus configures:

  - One or more addresses from PREFIX_A.

  - Two default routers (ROUTER_A and ROUTER_B).

  - One recursive DNS server: RDNSS_A.

Consider the case where e.g. ROUTER_A becomes unreachable, while ROUTER_B is still known to be reachable. If ROUTER_A is still in the reachable state:

* Ongoing communication instances currently employing ROUTER_A as their next-hop should switch to employing ROUTER_B.

* New communication instances should employ ROUTER_B as their next-hop (since ROUTER_B is known to be reachable, while ROUTER_A is known to be unreachable).

Existing implementations are expected to handle this scenario gracefully, since Section 6.3.6 of [RFC4861] required that "Routers that are reachable or probably reachable (i.e., in any state other than INCOMPLETE) SHOULD be preferred over routers whose reachability is unknown or suspect".

3.4. Multi-Interface Scenario

This scenario is similar to the one described in Section 3.1 with the only difference in that a host is attached to one or more networks via two or more network interfaces.

NOTE:
In the most common multi-interface case, the host will be attached to different networks via each network interfaces, and hence the multi-interface scenario will typically also result in a "multi-router" and "multiple-prefix" scenario (see Section 2). However, in the specific corner case where a host employs two network interfaces two attach to the same local network, this will essentially result in a Section 3.1 scenario (see Section 2): while the router available via both interfaces is actually the same router, from the perspective of the host it can be considered to be a different router (e.g., they will typically result in two default routers, each via the corresponding network interface).

Consider a network scenario where a user connects to two different ISPs (ISP_A and ISP_B), via two different network interfaces (e.g., one Ethernet interface and a wireless Wi-Fi interface). The scenario could be described as follows:

* ROUTER_A from ISP_A is attached to Network_A, and advertises:

  - Prefix PREFIX_A for address configuration (by means of a Prefix Information Option (PIO) [RFC4861]).

  - One Recursive DNS server (RDNSS_A) (by means of Recursive DNS Server option [RFC8106]).

* ISP_A enforces ingress filtering [RFC2827], and implements ACLs such that RDNSS_A only processes requests from ISP_A customers.

* ROUTER_B from ISP_B is attached to Network_B, and advertises:

  - Prefix PREFIX_B for address configuration (by means of a Prefix Information Option (PIO) [RFC4861]).

  - One Recursive DNS server (RDNSS_B) (by means of Recursive DNS Server option [RFC8106]).

* ISP_B enforces ingress filtering [RFC2827], and implements ACLs such that RDNSS_B only processes requests from ISP_B customers.

* Host C attaches to Network_A with one network interface, and to Network_B with another network interface, and configures:

  - Addresses from both prefixes (PREFIX_A and PREFIX_B).

  - Two default routers (ROUTER_A and ROUTER_B).

  - Two recursive DNS servers: RDNSS_A and RDNSS_B.

In this scenario, Host C may only send traffic from PREFIX_A via ROUTER_A or from PREFIX_B via ROUTER_B: otherwise, packets will be dropped as a result of ingress filtering [RFC2827]. Similarly, Host C may only send DNS queries from PREFIX_A to RDNSS_A, or from PREFIX_B to RDNSS_B: sending traffic from PREFIX_A to RDNSS_B or from PREFIX_B to RDNSS_A will resul in the ACLs enforced by the respective ISPs to drop the DNS queries.

It should be noted that it is quite common for DNS responses to depend on the source address of the query, for improved service. For example, if the "www.example.com" web site was served via a Content Delivery Network (CDN), RDNSS_A would likely resolve that domain name to an IP address that is topologically close to ISP_A, while RDNSS_B would resolve the same domain name to an address that is topologically close to ISP_B. In many cases, the corresponding web cache might be hosted within the ISP network itself. In these scenarios, using the information learned from RDNSS_A via ISP_B, or using the information learned via RDNSS_B via ISP_A would likely lead to sub-optimal service (e.g., larger Round-Trip Time) or no service (if ACLs were being enforced).

It should be evident that each piece of information being advertised via SLAAC is only usable when employed in conjunction with the rest of the information advertised by the same router. However, SLAAC does not require this behavior: hosts are free to use any piece of configuration they learn via SLAAC as they see fit.

NOTE:
For example, a host implementing the Weak End System (ES) model (see Section 3.3.4.2 from [RFC1122]) could indeed send e.g. packets from PREFIX_A via ROUTER_B.

3.5. Conflicting Information

Consider the case where two routers attach to the same network, and advertise the same configuration information. That is,

* Two SLAAC routers (ROUTER_A and ROUTER_B) are attached to Network_C

* ROUTER_A advertises:

  - Prefix PREFIX_A for address configuration (by means of a Prefix Information Option (PIO) [RFC4861]).

  - One Recursive DNS server (RDNSS_A) (by means of Recursive DNS Server option [RFC8106]).

* ROUTER_B advertises:

  - Prefix PREFIX_A for address configuration (by means of a Prefix Information Option (PIO) [RFC4861]).

  - One Recursive DNS server (RDNSS_A) (by means of Recursive DNS Server option [RFC8106]).

* Host C attaches to Network_C, and thus configures:

  - One or more addresses from PREFIX_A.

  - Two default routers (ROUTER_A and ROUTER_B).

  - One recursive DNS server: RDNSS_A.

Consider the case where e.g. ROUTER_B is unable to refresh its network configuration information from its upstream, and thus advertises the same configuration as before, but with a lifetime of 0. That is, it advertises:

* A PIO conveying PREFIX_A with both a Preferred Lifetime and a Valid Lifetime of 0.

* A RDNSS conveying RDNSS_A with a Lifetime of 0.

Presumably, this means that according to ROUTER_B, this information should no longer be used:

* Hosts should remove any configured addresses for such prefixes. As a result, they should also abort any ongoing TCP connections.

* Hosts should also remove the corresponding RDNSS server from their list of RDNSS servers.

This would also happen if ROUTER_A was still announcing the same configuration information with non-zero lifetimes.

It is clear that a more resilient behavior would be to maintain, for each piece of network configuration information, different timers for each SLAAC advertising router. Thus, if a SLAAC router advertised some configuration information with a lifetime of 0, this would simply mean that such configuration information should be disassociated with that particular router. Only when configuration information is no longer associated with any router would the information be removed from the host altogether.

4. Prior Work

[RFC8028] analyzes the challenge represented by having multiple default routers when addresses from multiple prefixes are employed. However, there are a few gaps in the specification:

* Extrapolating RFC 8028 to other network configuration information (such as Route Information Options (RIOs) [RFC4191] and RDNSS [RFC8106]), as discussed in Section 3.1 of this document.

* Considering how to aggregate configuration information when the same information is advertised by multiple routers, with different timers/lifetime values (as discussed in Section 3.5).

* Considering failure cases, such as those discussed in Section 3.3 and Section 3.2.

5. Future Work

This document describes a number of common network scenarios that are currently unsupported by IPv6. These scenarios have become more and more common, as a result of:

* Increased number of home-office users, requiring the use of multiple upstream ISP for improved resiliency

* Increased number of mobile users, which may not only connect via the mobile operator but also via a Wi-Fi connection when available.

As a result, this document concludes that protocol improvements that accommodate these deployment scenarios are warranted. [I-D.gont-6man-multi-ipv6-spec] is a draft protocol specification that aims to incorporate support for these scearios into IPv6 hosts.

6. IANA Considerations

This document has no actions for IANA.

7. Security Considerations

This document does not introduce any new attack vectors. A host that were to implement the behavior described in this document might actually reduce the impact of some Neighbor discovery attacks. For example, an ND attack meaning to disable an IPv6 prefix by forging a Router Advertisement (RA) with a PIO for the target prefix with a zero lifetime would only succeed if:

* the RA impersonates an existing router (i.e., it employs the address of an existing router as the source address of the RA packet).

* No other router on the same network segment is currently advertising the target prefix.

Similarly, in a scenario where a host is employing multiple interfaces, and an attacker tries to disable the usage of a RDNSS by sending forged RAs advertising a RDNSS with a zero lifetime, the attacker would only be able to affect usage of that RDNSS via the network interface attached to network on which the attacker is performing the attack. However, RDNSS servers employed via network interfaces attached to different networks would remain unaffected.

8. Acknowledgments

The authors would like to thank (in alphabetical order) Brian Carpenter for providing valuable comments on earlier 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 has benefited his protocol-related work.

9. References

9.1. Normative References

[RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989, <https://www.rfc-editor.org/info/rfc1122>.

[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>.

[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>.

[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>.

[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>.

9.2. Informative References

[I-D.gont-6man-multi-ipv6-spec]
Gont, F., "Improving Support for Multi-Router and Multi-Prefix IPv6 Networks", Work in Progress, Internet-Draft, draft-gont-6man-multi-ipv6-spec-00, 2 February 2025, <https://datatracker.ietf.org/doc/html/draft-gont-6man-multi-ipv6-spec-00>.

[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>.

Author’s Address

Fernando Gont
SI6 Networks
Segurola y Habana 4310, 7mo Piso
Villa Devoto
Ciudad Autonoma de Buenos Aires
Argentina
Email: fgont@si6networks.com
URI: https://www.si6networks.com