6man WG
Internet-Draft
Updates: 6724 (if approved)
Intended status: Standards Track
Expires: October 1, 2017

J. Linkova
Google

March 30, 2017

Default Address Selection and Subnet Renumbering
draft-linkova-6man-default-addr-selection-update-00

Abstract

This document discusses some scenarios when IPv6 hosts might not be able to properly detect the fact the network they are connected to has changed IPv6 addressing. It proposes changes to the Default Address Selection algorithm defined in [RFC6724] to mitigate the impact of the abovementioned failure scenarios as well as provides recommendations for sending Prefix Information Options (PIO). It updated [RFC6724].

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 http://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 October 1, 2017.

Table of Contents

1. Introduction

When an IPv6 host configures an address using Stateless Address Autoconfiguration (SLAAC) as described in [RFC4862], the configured address stays preferred (and therefore can be used for new communications) until one of the following happens:

- its preferred lifetime expires
- the hosts receives an router advertisement (RA) with the corresponding Prefix Information Option (PIO) with preferred lifetime set to zero
- the network interface changes its status

In other words once a host get connected to a network and an IPv6 address is configured that address may be used for quite long time (the default value of preferred lifetime is 7 days) or until the host received an explicit notification from a router that the particular SLAAC prefix is not valid anymore.

A host might need to stop using addresses from a particular prefix in the following scenarios:

- the host has moved to another layer 2 domain (e.g. VLAN or LAN)
- the layer 2 domain the host is connected to has been renumbered to another /64

In the ideal world the first scenario (a host moving to another layer 2 domain) would trigger the interface status change and as a result all network settings being reset. In the second scenario (network renumbering) it is expected that the router is sending an RA with the "old" PIO preferred lifetime set to zero and then a new POI is sent so hosts can use that POI for SLAAC. In either case the host receives an explicit notification about the addressing change. The preferred lifetime value is acting as a "safety net", with the default value being 604800 seconds (7 days) ([RFC4861]) and the realistic minimal value at least 12 seconds in the best case scenario being too long to rely on to detect the address change.

Unfortunately in practice there are some scenarios when a failure (or misconfiguration) on the host or the network level leds to a situation when a host is using addresses from a prefix which should e deprecated as it is not assigned to that layer 2 domain anymore. This results in a host using a "wrong" IPv6 address for initiating the connection and, as the returning packets can not reach the host, broken IPv6 connectivity and unsatisfactory user experience. Therefore it would be desirable to explore the feasibilty of updating hosts and routers behavior to minimize the impact and make IPv6 implementations more robust to such failures/misconfigurations.

1.1. 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 "Key words for use in RFCs to Indicate Requirement Levels" [RFC2119].

2. Failure Scenarios

Scenarios when a host might not receive an explicit notification leading to a prefix deprecation include but are not limited to:

- A switchport the host is connected to is moved to another subnet (VLAN) as a result of manual switchport reconfiguration or 802.1x re-authentication. In particular there have been evidence that some 802.1x supplicants do not reset network setting after successful 802.1x authentication. So if a host had failed 802.1x authentication for some reasons, was placed in a "quarantine" vlan and then got succesfully authenticated later on, it might end up having Ipv6 addresses from both old ("quarantine") and new vlans.

- A router which had received a prefix via DHCP-PD and sent RAs with the corresponding PIOs to hosts in LAN segments got rebooted/ crashed. After coming back up the router received a new DHCP-PD prefix so all connected hosts received RAs with a new POI.

- During the planned network renumbering a router was configured to send an RA with preferred lifetime for the "old" POI set to zero and the new POI having non-zero preferred lifetime. However due to unsolicited RAs being send as all-hosts multicast and the multicast being rather unreliable on busy wifi network, that RA was not received by a host

- Automated device config management system performs periodical config push to network devices. If such a push results in changing /64 subnet configured on a particular network, hosts attached to that network would not get notified about the subnet change and their addresses from the "old" prefix are not deprecated. The related case is incorrectly performed renumbering when a network administrator is renumbering a network by simply removing the "old" prefix from the configuration and configuring a new prefix instead.

All those (and others) scenarios result in a situation when the host has addresses from two different prefixes, "old" and "new". As both addresses are preferred and allowed to be used for communication the host relies on the default address selection algorithm ([RFC6724]) to choose a source address. If the address from the "old" prefix is selected as source address, then even if the packet reaches its destination (not being dropped due to antispoofing or any other type of filtering), the return traffic would not be delivered to the host.

3. Proposed Solition

3.1. Default Address Selection Algorithm Update

The Default Address Selection algorithm defined in [RFC6724] describes 8 rules to choose a single source address for use with a given IPv6 destination address. In the abovementioned scenario when the host has preferred addresses from two GUA prefixes, the first 7 rules can not act as a tie breaker. In theory when the host moves from one network segment from another its default router link-local address would change and the rule 5.5, "Prefer addresses in a prefix advertised by the next-hop" can lead to selecting a source address from the "new" prefix. However there are two reasons why the rule 5.5 can not reliable ensure that the "new" prefix is preferred over the "old" one:

1. The link-local address of the router in the new layer 2 domain might be the same as the link-local address of the "old" router (it's quite common to have link-local address on routers to be explicitly configured, especially in VRRP-enabled environments)

2. Until recently ([RFC8028]) IPv6 implementations were not required to track what next hop advertized what PIO and therefore the rule 5.5 was not applicable for such implementations.

The last rule, the rule 8, instructs the host to use the longest matching prefix and according to [RFC6724] that rule MAY be superseded if the implementation has other means of choosing among source addresses. In all scenarios described above it seems to be beneficial to prefer an address from the most recently received PIO. It would ensure that if the network subnet has been changed and the host has addresses from both "old" and "new" prefixes, it would prefer the new prefix. In generic case choosing an address from the most recent PIO if none of the first seven source address selection rules can be a tie breaker is harmless. If all POIs were received in the same time (the same RA) then the rule 8 (or any other means) can be used to choose the source address.

Therefore this document proposes the following changes to the Section 5 of [RFC6724]:

OLD TEXT:

Rule 8: Use longest matching prefix.

If CommonPrefixLen(SA, D) > CommonPrefixLen(SB, D), then prefer SA. Similarly, if CommonPrefixLen(SB, D) > CommonPrefixLen(SA, D), then prefer SB.

Rule 8 MAY be superseded if the implementation has other means of choosing among source addresses. For example, if the implementation somehow knows which source address will result in the "best" communications performance.

NEW TEXT:

Rule 8: Use the address from the most recently refreshed prefix.

If SA's PIO was received more recently than SB's POI, then prefer SA. Similarly, if SB's POI was received more recently than SA's POI, then prefer SB. If the implementation does not keep track of when the particular POI was received, than the addresses preferred lifetime SHOULD be considered instead: if preferred lifetime(SA) > preferred lifetime(SB), then prefer SA. Similarly, if preferred lifetime(SB) > preferred lifetime(SA), then prefer SB.

Rule 9: Use longest matching prefix.

If CommonPrefixLen(SA, D) > CommonPrefixLen(SB, D), then prefer SA. Similarly, if CommonPrefixLen(SB, D) > CommonPrefixLen(SA, D), then prefer SB.

Rules 8 and 9 MAY be superseded if the implementation has other means of choosing among source addresses. For example, if the implementation somehow knows which source address will result in the "best" communications performance.

To make the proposed solution work for the implementations which do not record when an RA with the PIO was most recently received, both old and new POI need to be advertised with same (or reasonably similar) preferred lifetime value. Otherwise it is possible that even the new POI was received after the old POI, the preferred lifetime of the old prefix might be still higher that one of the new prefix (if the preferred lifetime field value for the old prefix was much higher that the corresponding value for the new prefix). Despite such a limitation it seems reasonable to assume that in most scenarios described in Section 2 the PIOs preferred lifetime values would not vary much.

4. IANA Considerations

This memo asks the IANA for no new parameters.

5. Security Considerations

This memo has no direct security considerations.

6. Acknowledgements

The authors thank Erik Kline for input and contributions.

7. 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, <http://www.rfc-editor.org/info/rfc2119>.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <http://www.rfc-editor.org/info/rfc2460>.

[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, <http://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, <http://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, <http://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, <http://www.rfc-editor.org/info/rfc8028>.

Appendix A. Change Log

Initial Version: March 2017

Author's Address

Jen Linkova
Google
Mountain View, California 94043
USA

Email: furry@google.com