

Internet Engineering Task Force (IETF)
Request for Comments: 9131
Updates: 4861
Category: Standards Track
ISSN: 2070-1721

J. Linkova
Google
October 2021

Gratuitous Neighbor Discovery: Creating Neighbor Cache Entries on First-Hop Routers

Abstract

Neighbor Discovery (RFC 4861) is used by IPv6 nodes to determine the link-layer addresses of neighboring nodes as well as to discover and maintain reachability information. This document updates RFC 4861 to allow routers to proactively create a Neighbor Cache entry when a new IPv6 address is assigned to a node. It also updates RFC 4861 and recommends that nodes send unsolicited Neighbor Advertisements upon assigning a new IPv6 address. These changes will minimize the delay and packet loss when a node initiates connections to an off-link destination from a new IPv6 address.

Status of This Memo

This is an Internet Standards Track document.

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

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9131>.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction
 - 1.1. Requirements Language
 - 1.2. Terminology
2. Problem Statement
3. Solution Requirements
4. Changes to Neighbor Discovery
 - 4.1. Nodes Sending Gratuitous Neighbor Advertisements
 - 4.2. Routers Creating Cache Entries upon Receiving Unsolicited Neighbor Advertisements

- 5. Avoiding Disruption
 - 5.1. Neighbor Cache Entry Exists in Any State Other Than INCOMPLETE
 - 5.2. Neighbor Cache Entry Is in INCOMPLETE State
 - 5.3. Neighbor Cache Entry Does Not Exist
 - 5.3.1. The Rightful Owner Is Not Sending Packets from the Address
 - 5.3.2. The Rightful Owner Has Started Sending Packets from the Address
- 6. Modifications to RFC-Mandated Behavior
 - 6.1. Modification to RFC 4861 (Neighbor Discovery for IP version 6 (IPv6))
 - 6.1.1. Modification to Section 7.2.5 of RFC 4861
 - 6.1.2. Modification to Section 7.2.6 of RFC 4861
- 7. Solution Limitations
- 8. Solutions Considered but Discarded
 - 8.1. Do Nothing
 - 8.2. Change to the Registration-Based Neighbor Discovery
 - 8.3. Host Sending NS to the Router Address from Its GUA
 - 8.4. Host Sending Router Solicitation from Its GUA
 - 8.5. Routers Populating Their Caches by Gleaning from Neighbor Discovery Packets
 - 8.6. Initiating Host-to-Router Communication
 - 8.7. Making the Probing Logic on Hosts More Robust
 - 8.8. Increasing the Buffer Size on Routers
 - 8.9. Transit Data Plane Traffic from a New Address to Trigger Address Resolution
- 9. IANA Considerations
- 10. Security Considerations
- 11. References
 - 11.1. Normative References
 - 11.2. Informative References
- Acknowledgements
- Author's Address

1. Introduction

The Neighbor Discovery state machine defined in [RFC4861] assumes that communications between IPv6 nodes are, in most cases, bidirectional and if a node A is trying to communicate to its neighbor, node B, the return traffic flows could be expected. So, when node A starts the address resolution process, the target node B would also create an entry containing A's IPv6 and link-layer addresses in its Neighbor Cache. That entry will be used for sending the return traffic to A.

In particular, Section 7.2.5 of [RFC4861] states:

```
| When a valid Neighbor Advertisement is received (either solicited
| or unsolicited), the Neighbor Cache is searched for the target's
| entry. If no entry exists, the advertisement SHOULD be silently
| discarded. There is no need to create an entry if none exists,
| since the recipient has apparently not initiated any communication
| with the target.
```

While this approach is perfectly suitable for host-to-host on-link communications, it does not work so well when a host sends traffic to off-link destinations. After joining the network and receiving a Router Advertisement, the host populates its Neighbor Cache with the default router IPv6 and link-layer addresses and is able to send traffic to off-link destinations. At the same time, the router does not have any cache entries for the host global addresses yet and only starts address resolution upon receiving the first packet of the return traffic flow. While waiting for the resolution to complete, routers only keep a very small number of packets in the queue, as recommended in Section 7.2.2 of [RFC4861]. Any additional packets

arriving before the resolution process finishes are likely to result in dropped packets. It can cause packet loss and performance degradation that can be visible to users.

This document updates the Neighbor Discovery protocol [RFC4861] to avoid packet loss in the scenario described above. Section 4 discusses the changes and analyzes the potential impact, while normative changes to [RFC4861] are specified in Section 6.

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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

1.2. Terminology

Node: A device that implements IP [RFC4861].

Host: Any node that is not a router [RFC4861].

ND: Neighbor Discovery [RFC4861].

NC: Neighbor Cache [RFC4861]. The Neighbor Cache entry can be in one of five states, as described in Section 7.3.2 of [RFC4861]: INCOMPLETE, REACHABLE, STALE, DELAY, or PROBE.

SLAAC: IPv6 Stateless Address Autoconfiguration [RFC4862].

NS: Neighbor Solicitation [RFC4861].

NA: Neighbor Advertisement [RFC4861].

RS: Router Solicitation [RFC4861].

RA: Router Advertisement [RFC4861].

SLLAO: Source Link-Layer Address Option. An option in the ND packets containing the link-layer address of the sender of the packet [RFC4861].

TLLAO: Target Link-Layer Address Option. An option in the ND packets containing the link-layer address of the target [RFC4861].

GUA: Global Unicast Address [RFC4291].

DAD: Duplicate Address Detection [RFC4862].

Preferred Address: An address assigned to an interface whose uniqueness has been verified using DAD and whose use by upper-layer protocols is unrestricted [RFC4862]. Preferred addresses may be used as the source address of packets sent from the interface.

Optimistic DAD: A modification of DAD [RFC4429].

2. Problem Statement

The most typical scenario when the problem described in this document may arise is a host joining the network, forming a new address, and using that address for accessing the Internet:

1. A host joins the network and receives a Router Advertisement (RA) packet from the first-hop router (either a periodic unsolicited

RA or a response to a Router Solicitation sent by the host). The RA contains information the host needs to perform SLAAC and to configure its network stack. The RA is sent from the router's link-local address to a link-local destination address and may contain the link-layer address of the router. As a result, the host can populate its Neighbor Cache with the router's link-local and link-layer addresses.

2. The host starts opening connections to off-link destinations. A very common use case is a mobile device sending probes to detect Internet connectivity and/or the presence of a captive portal on the network. To speed up that process, many implementations use Optimistic DAD, which allows them to send probes before the DAD process is completed. At that moment, the device's Neighbor Cache contains all information required to send those probes (such as the default router link-local and link-layer addresses). The router's Neighbor Cache, however, might contain an entry for the device's link-local address (if the device has been performing address resolution for the router's link-local address), but there are no entries for any of the device's global addresses.
3. Return traffic is received by the first-hop router. As the router does not have any cache entry for the host's global address yet, the router starts the Neighbor Discovery process by creating an INCOMPLETE cache entry and then sending a Neighbor Solicitation to the solicited-node multicast address (Section 7.3.2 of [RFC4861]). As per Section 7.2.2 of [RFC4861], routers MUST buffer at least one data packet and MAY buffer more, while resolving the packet destination address. However, most router implementations limit the buffer size to a few packets only, and some implementations are known to buffer just one packet. So, any subsequent packets arriving before the address resolution process is completed cause packet loss by replacing older packets in the buffer.
4. If the host sends multiple probes in parallel, in the worst case, it would consider all but one of them failed. That leads to user-visible delay in connecting to the network, especially if the host implements some form of backoff mechanism and does not retransmit the probes as soon as possible.

This scenario illustrates the problem occurring when the device connects to the network for the first time or after an inactivity period long enough for the device's address to be removed from the router's Neighbor Cache. However, the same sequence of events happens when the host starts using a new global address previously unseen by the router, such as a new privacy address [RFC8981] or if the router's Neighbor Cache has been flushed.

While in dual-stack networks this problem might be hidden by Happy Eyeballs [RFC8305], it manifests quite clearly in IPv6-only environments, especially wireless environments, leading to poor user experience and contributing to a negative perception of IPv6-only solutions as unstable and non-deployable.

3. Solution Requirements

It would be highly desirable to improve the Neighbor Discovery mechanics so routers have a usable cache entry for a host address by the time the router receives the first packet for that address. In particular:

- * If the router does not have a Neighbor Cache entry for the address, a STALE entry needs to be created proactively, prior to arrival of the first packet intended for that address.

- * The solution needs to work for Optimistic Addresses as well. Devices implementing Optimistic DAD usually attempt to minimize the delay in connecting to the network and therefore are more likely to be affected by the problem described in this document.
- * In the case of duplicate addresses present in the network, the solution should not override the existing entry.
- * In topologies with multiple first-hop routers, the cache needs to be updated on all of them, as traffic might be asymmetric: outgoing flows leaving the network via one router while the return traffic enters the segment via another one.

In addition, the solution must not exacerbate issues described in [RFC6583] and needs to be compatible with the recommendations provided in [RFC6583].

4. Changes to Neighbor Discovery

The following changes are required to minimize the delay in creating new entries in a router's Neighbor Cache:

- * A node sends unsolicited NAs upon assigning a new IPv6 address to its interface.
- * A router creates a new cache entry upon receiving an unsolicited NA from a host.

The following sections discuss these changes in more detail. Normative changes are specified in Section 6.

4.1. Nodes Sending Gratuitous Neighbor Advertisements

Section 7.2.6 of [RFC4861] discusses using unsolicited Neighbor Advertisements to inform node neighbors of the new link-layer address quickly. The same mechanism could be used to notify the node neighbors about the new network-layer address as well: the node can send unsolicited Neighbor Advertisements upon assigning a new IPv6 address to its interface.

To minimize potential disruption in the case of duplicate addresses, the node should not set the Override flag for a preferred address and must not set the Override flag if the address is in the Optimistic state [RFC4429].

As the main purpose of sending unsolicited NAs upon configuring a new address is to proactively create a Neighbor Cache entry on the first-hop routers, the gratuitous NAs are sent to the all-routers multicast address (ff02::2). Limiting the recipients to routers only would help reduce the multicast noise level. If the link-layer devices are performing Multicast Listener Discovery (MLD) snooping [RFC4541], then those unsolicited NAs will only be sent to routers on the given network segment/link, instead of being flooded to all nodes.

It should be noted that the mechanism discussed here does not cause any significant increase in multicast traffic. The additional multicast unsolicited NAs would proactively create a STALE cache entry on the router, as discussed below. When the router receives the return traffic flows, it does not need to send multicast NSes to the solicited-node multicast address but would send unicast NSes instead. Therefore, this procedure would only produce an increase in the overall amount of multicast traffic if no return traffic arrives for the address that sent the unsolicited NA or if the router does not create a STALE entry upon receiving such an NA. The increase would be negligible, as that additional traffic is a few orders of

magnitude less than the usual level of Neighbor Discovery multicast traffic.

4.2. Routers Creating Cache Entries upon Receiving Unsolicited Neighbor Advertisements

Section 7.2.5 of [RFC4861] states:

```
| When a valid Neighbor Advertisement is received (either solicited  
| or unsolicited), the Neighbor Cache is searched for the target's  
| entry. If no entry exists, the advertisement SHOULD be silently  
| discarded. There is no need to create an entry if none exists,  
| since the recipient has apparently not initiated any communication  
| with the target.
```

The reasoning behind dropping unsolicited Neighbor Advertisements ("the recipient has apparently not initiated any communication with the target") is valid for on-link host-to-host communication but, as discussed in Section 1, it does not really apply to the scenario when the host is announcing its address to routers. Therefore, it would be beneficial to allow routers to create new entries upon receiving an unsolicited Neighbor Advertisement.

This document updates [RFC4861] so that routers create a new Neighbor Cache entry upon receiving an unsolicited Neighbor Advertisement for an address that does not already have a Neighbor Cache entry. These changes do not modify the router behavior specified in [RFC4861] for the scenario when the corresponding Neighbor Cache entry already exists.

The next section analyzes various scenarios of duplicate addresses and discusses the potential impact of creating a STALE entry for a duplicate IPv6 address.

5. Avoiding Disruption

If nodes following the recommendations in this document are using the DAD mechanism defined in [RFC4862], they would send unsolicited NAs as soon as the address changes state from tentative to preferred (after its uniqueness has been verified). However, nodes willing to minimize network stack configuration delays might be using Optimistic Addresses, which means there is a possibility of the address not being unique on the link. Section 2.2 of [RFC4429] discusses measures to ensure that ND packets from the Optimistic Address do not override any existing Neighbor Cache entries, as it would cause interruption of the rightful address owner's traffic in the case of an address conflict. Nodes that are willing to speed up their network stack configuration are most likely to be affected by the problem outlined in this document; therefore, it seems reasonable for such hosts to advertise their Optimistic Addresses by sending unsolicited NAs. The main question to consider is the potential risk of overriding the cache entry for the rightful address owner if the Optimistic Address happens to be a duplicate.

The following sections discuss the address collision scenario when a node sends an unsolicited NA for an address in the Optimistic state, while another node (the rightful owner) already has the same address assigned. This document uses the term "the rightful owner", as the same terminology is used in [RFC4429]. The analysis assumes that the host performs DAD, as Section 5.4 of [RFC4862] requires that DAD MUST be performed on all unicast addresses prior to assigning them to an interface.

5.1. Neighbor Cache Entry Exists in Any State Other Than INCOMPLETE

If the router's Neighbor Cache entry for the target address already

exists in any state other than INCOMPLETE, then as per Section 7.2.5 of [RFC4861], an unsolicited NA with the Override flag cleared would change the entry state from REACHABLE to STALE but would not update the entry in any other way. Therefore, even if the host sends an unsolicited NA from its Optimistic Address, the router's cache entry would not be updated with the new link-layer address, and no impact on the traffic for the rightful address owner is expected.

The return traffic intended for the host with the Optimistic Address would be sent to the rightful owner. However, this is unavoidable with or without the unsolicited NA mechanism.

5.2. Neighbor Cache Entry Is in INCOMPLETE State

Another corner case is the INCOMPLETE cache entry for the address.

1. The router receives a packet for the rightful owner of the address.
2. The router starts the address resolution process by creating an INCOMPLETE entry and sends the multicast NS.
3. More packets arrive at the router for the address in question.
4. The host configures an Optimistic Address and sends an unsolicited NA.
5. The router creates a STALE entry and sends the buffered packet(s) to the host (while at least some of those packets are actually intended for the rightful owner).
6. As the STALE entry was used to send packets, the router changes the entry state to DELAY and waits up to DELAY_FIRST_PROBE_TIME (5 seconds) [RFC4861] before sending a unicast NS.
7. The rightful owner responds to the multicast NS sent at Step 2 with a solicited NA with the Override flag set.
8. The router updates the entry with the TLLAO supplied (the rightful owner's link-layer address) and sets the entry state to REACHABLE (as the NA has the Solicited flag set).

As a result, some packets (packets in the buffer at Step 6 and all packets arriving between Step 6 and Step 8) are delivered to the host with the Optimistic Address, while some of them, if not all, are intended for the rightful owner. Without the unsolicited NA, one or more packets that are in the buffer at Step 8 (usually just one packet, but some routers may buffer a few) would have been delivered to the rightful owner and the rest of the packets would have been dropped. However, the probability of such a scenario is rather low, as it would require the following things to happen almost simultaneously (within tens of milliseconds in most cases):

- * One host starts using a new IPv6 address and sending traffic without sending an unsolicited NA first.
- * Another host configures the same IPv6 address in Optimistic mode before the router completes the address resolution process for the rightful owner.

It should be noted that in this scenario the rightful owner does not send any unsolicited NAs before sending packets. If the rightful owner implements the functionality described in this document and sends unsolicited NAs upon configuring its address, then the router creates a STALE entry for the address, causing all packets to be delivered to the rightful owner (see Section 5.1). The rightful

owner would experience no disruption but might receive some packets intended for the host with an Optimistic Address.

This section focuses on the scenario when the solicited NA from the rightful owner arrives after the unsolicited one sent from the Optimistic Address (Step 7 and Step 4, respectively). If the solicited NA arrives first, it changes the NC entry state from INCOMPLETE to REACHABLE. As discussed in Section 5.1, there will be no disruption for the rightful owner if the router already has a REACHABLE entry for the address when an unsolicited NA is received.

5.3. Neighbor Cache Entry Does Not Exist

There are two distinct scenarios that can lead to the situation when the router does not have an NC entry for the IPv6 address:

1. The rightful owner of the address has not been using it for off-link communication recently or has never used it at all.
2. The rightful owner just started sending packets from that address, but the router has not received any return traffic yet.

The impact on the rightful owner's traffic flows would be different in those cases.

5.3.1. The Rightful Owner Is Not Sending Packets from the Address

In this scenario, the following events are expected to happen:

1. The host configures the address and sets its state to Optimistic.
2. The host sends an unsolicited NA with the Override flag set to zero and starts sending traffic from the Optimistic Address.
3. The router creates a STALE entry for the address and the host link-layer address.
4. The host starts DAD and detects the address duplication.
5. The router receives the return traffic for the duplicate address. As the NC entry is STALE, it sends traffic using that entry, changes it to DELAY, and waits up to DELAY_FIRST_PROBE_TIME seconds [RFC4861].
6. The router changes the NC entry state to PROBE and sends up to MAX_UNICAST_SOLICIT unicast NSes [RFC4861] separated by RetransTimer milliseconds [RFC4861] to the host link-layer address.
7. As the host has already detected the address conflict, it does not respond to the unicast NSes. (It is unlikely that the host has not completed the DAD process at this stage, as DELAY_FIRST_PROBE_TIME (5 seconds) is much higher than the DAD duration ($\text{DupAddrDetectTransmits} \times \text{RetransTimer} \times 1000 + \text{MAX_RTR_SOLICITATION_DELAY}$ seconds) (Section 5.4 of [RFC4862]).) The default value for the DAD process would be $1 \times 1 \times 1000 + 1 = 2$ seconds [RFC4861]. If the host has completed DAD but did not detect the address conflict, then there are two hosts with the same address in the preferred state and disruption is inevitable anyway.
8. As the router receives no response for the unicast NSes, it deletes the NC entry.
9. If return packets for communication initiated at Step 2 are still arriving, the router buffers a small number of those packets and

starts the address resolution process again by sending a multicast NS to the solicited-node multicast address. The rightful owner responds, and the router's NC entry is updated with the rightful owner's link-local address. The buffered packet or packets are sent to that address. Any packets still arriving after the address resolution process has completed are sent to the rightful address owner as well.

The rightful owner is not experiencing any disruption, as it does not send any traffic. It would only start receiving packets intended for another host after Step 8 is completed and only if return packets for the communication initiated at Step 2 are still arriving.

However, the same behavior would be observed if the changes specified in this document are not implemented. If the host starts sending packets from its Optimistic Address but then detects that the address is a duplicate, the first return packet would trigger the address resolution process and would be buffered until the resolution is completed. The buffered packet(s) and any packets still arriving after the address is resolved would be forwarded to the rightful owner of the address. So, the rightful owner might still receive one or more packets from the flows intended for another host. Therefore, it's safe to conclude that the changes specified in this document do not introduce any disruption for the rightful owner of the duplicated address.

5.3.2. The Rightful Owner Has Started Sending Packets from the Address

In this scenario, the following events are happening:

1. The rightful owner starts sending traffic from the address (e.g., the address has just been configured or has not been recently used).
2. The host configures the address and sets its state to Optimistic.
3. The host sends an unsolicited NA with the Override flag set to zero and starts sending traffic from the Optimistic Address.
4. The router creates a STALE entry for the address and the host link-layer address.
5. The host starts DAD and detects the address duplication.
6. The router receives the return traffic for the IPv6 address in question. Some flows are intended for the rightful owner of the duplicate address, while some are for the new host. As the NC entry is STALE, it sends traffic using that entry, changes it to DELAY, and waits up to DELAY_FIRST_PROBE_TIME seconds [RFC4861].
7. The router changes the NC entry state to PROBE and sends up to MAX_UNICAST_SOLICIT unicast NSes [RFC4861] separated by RetransTimer milliseconds [RFC4861] to the host link-layer address.
8. As the host has already detected the address conflict, it does not respond to the unicast NSes.
9. As the router receives no response for the unicast NSes, it deletes the NC entry.
10. The next packet recreates the entry and triggers the resolution process. The router buffers the packet and sends a multicast NS to the solicited-node multicast address. The rightful owner responds, and the router's NC entry is updated with the rightful

owner's link-local address.

As a result, the traffic for the address of the rightful owner would be sent to the host with the duplicate address instead. The duration of the disruption can be estimated as $\text{DELAY_FIRST_PROBE_TIME} \times 1000 + (\text{MAX_UNICAST_SOLICIT} - 1) \times \text{RetransTimer}$ milliseconds. As per the constants defined in Section 10 of [RFC4861], this interval is equal to $5 \times 1000 + (3 - 1) \times 1000 = 7000$ milliseconds, or 7 seconds.

However, it should be noted that the probability of such a scenario is rather low. Similar to the scenario discussed in Section 5.2, it would require the following things to happen almost simultaneously (within tens of milliseconds in most cases):

- * One host starts using a new IPv6 address and sending traffic without sending an unsolicited NA first.
- * Another host configures the same IPv6 address in Optimistic mode before the router receives the return traffic for the first host.

As discussed in Section 5.2, the disruption for the rightful owner can easily be prevented if that node implements the mechanism described in this document. Sending unsolicited NAs before initiating off-link communication would create a STALE entry in the router's NC and prevent any traffic to that address from being sent to the host with the Optimistic Address (see Section 5.1).

6. Modifications to RFC-Mandated Behavior

All normative text in this memo is contained in this section.

6.1. Modification to RFC 4861 (Neighbor Discovery for IP version 6 (IPv6))

6.1.1. Modification to Section 7.2.5 of RFC 4861

This document makes the following changes to Section 7.2.5 of [RFC4861]:

The text in RFC 4861 is as follows:

When a valid Neighbor Advertisement is received (either solicited or unsolicited), the Neighbor Cache is searched for the target's entry. If no entry exists, the advertisement SHOULD be silently discarded. There is no need to create an entry if none exists, since the recipient has apparently not initiated any communication with the target.

This document updates the text as follows:

When a valid Neighbor Advertisement is received (either solicited or unsolicited), the Neighbor Cache is searched for the target's entry. If no entry exists:

- * Hosts SHOULD silently discard the advertisement. There is no need to create an entry if none exists, since the recipient has apparently not initiated any communication with the target.
- * Routers SHOULD create a new entry for the target address with the link-layer address set to the Target Link-Layer Address Option (if supplied). The entry's reachability state MUST be set to STALE. If the received Neighbor Advertisement does not contain the Target Link-Layer Address Option, the advertisement SHOULD be silently discarded.

6.1.2. Modification to Section 7.2.6 of RFC 4861

This document makes the following changes to Section 7.2.6 of [RFC4861]:

The text in RFC 4861 is as follows:

Also, a node belonging to an anycast address MAY multicast unsolicited Neighbor Advertisements for the anycast address when the node's link-layer address changes.

This document updates the text as follows:

Also, a node belonging to an anycast address MAY multicast unsolicited Neighbor Advertisements for the anycast address when the node's link-layer address changes.

A node may also wish to notify its first-hop routers when it configures a new global IPv6 address so the routers can proactively populate their Neighbor Caches with the corresponding entries. In such cases, a node SHOULD send up to MAX_NEIGHBOR_ADVERTISEMENT Neighbor Advertisement messages. If the address is preferred, then the Override flag SHOULD NOT be set. If the address is in the Optimistic state, then the Override flag MUST NOT be set. The destination address SHOULD be set to the all-routers multicast address. These advertisements MUST be separated by at least RetransTimer seconds. The first advertisement SHOULD be sent as soon as one of the following events happens:

If Optimistic DAD [RFC4429] is used: A new Optimistic Address is assigned to the node interface.

If Optimistic DAD is not used: An address changes the state from tentative to preferred.

7. Solution Limitations

The solution described in this document provides some improvement for a node configuring a new IPv6 address and starting to send traffic from it. However, that approach does not completely eliminate the scenario when a router receives some transit traffic for an address without the corresponding Neighbor Cache entry. For example:

- * If the host starts using an already-configured IPv6 address after a long period of inactivity, the router might not have the NC entry for that address anymore, as old/expired entries are deleted.
- * Clearing the router's Neighbor Cache would trigger packet loss for all actively used addresses removed from the cache.

8. Solutions Considered but Discarded

There are other possible approaches to address the problem. For example:

- * Just do nothing.
- * Migrate from the "reactive" Neighbor Discovery [RFC4861] to the registration-based mechanisms [RFC8505].
- * Create new entries in the router's Neighbor Cache by gleaning from Neighbor Discovery DAD messages.
- * Initiate bidirectional communication from the host to the router using the host GUA.

- * Make the probing logic on hosts more robust.
- * Increase the buffer size on routers.
- * Transit data plane traffic from an unknown address (an address without the corresponding Neighbor Cache entry) to trigger an address resolution process on the router.

It should be noted that some of those options are already implemented by some vendors. The following sections discuss those approaches and the reasons they were discarded.

8.1. Do Nothing

One of the possible approaches might be to declare that everything is working as intended and let the upper-layer protocols deal with packet loss. The obvious drawbacks include:

- * Unhappy users.
- * Many support tickets.
- * More resistance to deploying IPv6 and IPv6-only networks.

8.2. Change to the Registration-Based Neighbor Discovery

The most radical approach would be to move away from the reactive ND as defined in [RFC4861] and expand the registration-based ND [RFC6775] [RFC8505] used in IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs) to the rest of the IPv6 deployments. This option requires some investigation and discussion. However, significant changes to the existing IPv6 implementations would be needed, so an unclear adoption timeline makes this approach less preferable than the approach specified in this document.

8.3. Host Sending NS to the Router Address from Its GUA

The host could force the creation of a STALE entry for its GUA in the router's Neighbor Cache by sending the following Neighbor Solicitation message:

- * The NS source address is the host GUA.
- * The destination address is the default router IPv6 address.
- * The Source Link-Layer Address Option contains the host link-layer address.
- * The target address is the host's default router address (the default router address the host received in the RA).

The main disadvantages of this approach are as follows:

- * It would not work for Optimistic Addresses, as Section 2.2 of [RFC4429] explicitly prohibits sending Neighbor Solicitations from an Optimistic Address.
- * If first-hop redundancy is deployed in the network, the NS would reach the active router only, so all backup routers (or all active routers except one) would not get their Neighbor Cache updated.
- * Some wireless devices are known to alter ND packets and perform various nonobvious forms of ND proxy actions. In some cases, unsolicited NAs might not even reach the routers.

8.4. Host Sending Router Solicitation from Its GUA

The host could send a Router Solicitation message to the all-routers multicast address, using its GUA as a source. If the host link-layer address is included in the Source Link-Layer Address Option, the router would create a STALE entry for the host GUA as per Section 6.2.6 of [RFC4861]. However, this approach cannot be used if the GUA is in the Optimistic state: Section 2.2 of [RFC4429] explicitly prohibits using an Optimistic Address as the source address of a Router Solicitation with a SLLAO, as it might cause disruption for the rightful owner of the address in the case of a collision. So, for the Optimistic Addresses, the host can send an RS without a SLLAO included. In that case, the router may respond with either a multicast or unicast RA (only the latter would create a cache entry).

This approach has the following drawbacks:

- * If the address is in the Optimistic state, the RS cannot contain a SLLAO. As a result, the router would only create a cache entry if solicited RAs are sent as unicast. Routers sending solicited RAs as multicast would not create a new cache entry, as they do not need to send a unicast packet back to the host.
- * There might be a random delay between receiving an RS and sending a unicast RA back (and creating a cache entry), which might undermine the idea of creating the cache entry proactively.
- * Some wireless devices are known to intercept ND packets and perform various nonobvious forms of ND proxy actions. In some cases, the RS might not even reach the routers.

8.5. Routers Populating Their Caches by Gleaning from Neighbor Discovery Packets

Routers may be able to learn about new addresses by gleaning from the DAD Neighbor Solicitation messages. The router could listen to all solicited-node multicast address groups and, upon receiving a Neighbor Solicitation from the unspecified address, search its Neighbor Cache for the solicitation's target address. If no entry exists, the router may create an entry, set its reachability state to INCOMPLETE, and start the address resolution process for that entry.

The same solution was proposed in [ND-ADDR-RES]. Some routing vendors already support such optimization. However, this approach has a number of drawbacks and therefore should not be used as the only solution:

- * Routers need to receive all multicast Neighbor Discovery packets; this might negatively impact a router's CPU.
- * If the router starts the address resolution process as soon as it receives the DAD Neighbor Solicitation, the host might still be performing DAD and the target address might be tentative. In that case, the host SHOULD silently ignore the received Neighbor Solicitation from the router as per Section 5.4.3 of [RFC4862]. As a result, the router might not be able to complete the address resolution process before the return traffic arrives.

8.6. Initiating Host-to-Router Communication

The host may force the router to start address resolution by sending a data packet such as ping or traceroute to its default router link-local address, using the GUA as a source address. As the RTT to the default router is lower than the RTT to any off-link destinations, it's quite likely that the router would start the Neighbor Discovery process for the host GUA before the first packet of the returning

traffic arrives.

This approach has the following drawbacks:

- * Data packets to the router's link-local address could be blocked by a security policy or control plane protection mechanism.
- * It introduces an additional overhead for the router's control plane (in addition to processing ND packets, the data packet needs to be processed as well).
- * Unless the data packet is sent to the all-routers ff02::2 multicast address, if the network provides a first-hop redundancy, then only the active router would create a new cache entry.

8.7. Making the Probing Logic on Hosts More Robust

Theoretically, the probing logic on hosts might be modified to better deal with initial packet loss. For example, only one probe can be sent, or probe retransmit intervals can be reduced. However, this approach has a number of drawbacks:

- * It would require updating all possible applications that perform probing, while the solution described in this document is implemented at the operating-system level.
- * Some implementations need to send multiple probes. Examples include but are not limited to:
 - Sending AAAA and A record DNS probes in parallel.
 - Detecting captive portals, which often requires sending multiple packets.
- * While it would increase the probability that the probing will complete successfully, there are multiple cases when packet loss would still occur:
 - The probe response consists of multiple packets, so all but the first one are dropped.
 - There are multiple applications on the same host sending traffic, and return packets arrive simultaneously.
 - There are multiple first-hop routers in the network. The first probe packet creates the NC entry on one of them. The subsequent return traffic flows might cross other routers and still experience the issue.
- * Reducing the probe retransmit interval unnecessarily increases network utilization and might cause network congestion.

8.8. Increasing the Buffer Size on Routers

Increasing the buffer size and buffering more packets would exacerbate issues described in [RFC6583] and make the router more vulnerable to ND-based denial-of-service attacks.

8.9. Transit Data Plane Traffic from a New Address to Trigger Address Resolution

When a router receives a transit packet sourced by an on-link neighbor node, it might check for the presence of a Neighbor Cache entry for the packet source address and, if the entry does not exist, start the address resolution process. This approach does ensure that a Neighbor Cache entry is proactively created every time a new,

previously unseen GUA is used for sending off-link traffic. However, this approach has a number of limitations. In particular:

- * If traffic flows are asymmetrical, the return traffic might not transit the same router as the original traffic that triggered the address resolution process. So, the Neighbor Cache entry is created on the "wrong" router, not the one that actually needs the Neighbor Cache entry for the host address.
- * The functionality needs to be limited to explicitly configured networks/interfaces, as the router needs to distinguish between on-link addresses (addresses for which the router needs to have Neighbor Cache entries) and the rest of the address space. The proactive address resolution process must only be triggered by packets from the prefixes known to be on-link. Otherwise, traffic from spoofed source addresses or any transit traffic could lead to Neighbor Cache exhaustion.
- * Implementing such functionality is much more complicated than all other solutions, as it would involve complex interactions between the data plane and the control plane.

9. IANA Considerations

This document has no IANA actions.

10. Security Considerations

One of the potential attack vectors to consider is cache spoofing, where the attacker might try to install a cache entry for the victim's IPv6 address and the attacker's link-layer address. However, it should be noted that this document does not propose any changes for the scenario when the Neighbor Cache for a given IPv6 address already exists. Therefore, there are no new vectors for an attacker to override an existing cache entry.

Section 5 describes some corner cases when a host with a duplicate Optimistic Address might get some packets intended for the rightful owner of the address. However, such scenarios do not introduce any new attack vectors: even without the changes discussed in this document, an attacker can easily override the router's Neighbor Cache and redirect the traffic by sending NAs with the Solicited flag set. As discussed in Section 5.3.2, the worst-case scenario might cause a disruption for up to 7 seconds. Because this scenario is highly unlikely, this risk of disruption is considered acceptable. More importantly, for all cases described in Section 5, the rightful owner can prevent disruption caused by an accidental address duplication just by implementing the mechanism described in this document. If the rightful owner sends unsolicited NAs before using the address, the STALE entry would be created on the router's NC, and any subsequent unsolicited NAs sent from the host with an Optimistic Address would not override the NC entry.

A malicious host could attempt to exhaust the Neighbor Cache on the router by creating a large number of STALE entries. However, this attack vector is not new, and the mechanism specified in this document does not increase the risk of such an attack: the attacker could do it, for example, by sending an NS or RS packet with a SLLAO included. All recommendations from [RFC6583] still apply.

Announcing a new address to the all-routers multicast address may inform an on-link attacker about IPv6 addresses assigned to the host. However, hiding information about the specific IPv6 address should not be considered a security measure, as such information is usually disclosed via DAD to all nodes anyway if MLD snooping is not enabled. Network administrators can also mitigate this issue by enabling MLD

snooping on the link-layer devices to prevent IPv6 link-local multicast packets from being flooded to all on-link nodes. If peer-to-peer on-link communications are not desirable for a given network segment, they should be prevented by proper Layer 2 security mechanisms. Therefore, the risk of allowing hosts to send unsolicited Neighbor Advertisements to the all-routers multicast address is low.

It should be noted that the mechanism discussed in this document allows hosts to proactively inform their routers about global IPv6 addresses existing on-link. Routers could use that information to distinguish between used and unused addresses to mitigate Neighbor Cache exhaustion DoS attacks as described in Section 4.3.2 of [RFC3756] and in [RFC6583].

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>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.
- [RFC4429] Moore, N., "Optimistic Duplicate Address Detection (DAD) for IPv6", RFC 4429, DOI 10.17487/RFC4429, April 2006, <<https://www.rfc-editor.org/info/rfc4429>>.
- [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>>.
- [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>>.

11.2. Informative References

- [ND-ADDR-RES] Chen, I. and J. Halpern, "Triggering ND Address Resolution on Receiving DAD-NS", Work in Progress, Internet-Draft, draft-halpern-6man-nd-pre-resolve-addr-00, 10 January 2014, <<https://datatracker.ietf.org/doc/html/draft-halpern-6man-nd-pre-resolve-addr-00>>.
- [RFC3756] Nikander, P., Ed., Kempf, J., and E. Nordmark, "IPv6 Neighbor Discovery (ND) Trust Models and Threats", RFC 3756, DOI 10.17487/RFC3756, May 2004, <<https://www.rfc-editor.org/info/rfc3756>>.
- [RFC4541] Christensen, M., Kimball, K., and F. Solensky, "Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches", RFC 4541, DOI 10.17487/RFC4541, May 2006, <<https://www.rfc-editor.org/info/rfc4541>>.

- [RFC6583] Gashinsky, I., Jaeggli, J., and W. Kumari, "Operational Neighbor Discovery Problems", RFC 6583, DOI 10.17487/RFC6583, March 2012, <<https://www.rfc-editor.org/info/rfc6583>>.
- [RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <<https://www.rfc-editor.org/info/rfc6775>>.
- [RFC8305] Schinazi, D. and T. Pauly, "Happy Eyeballs Version 2: Better Connectivity Using Concurrency", RFC 8305, DOI 10.17487/RFC8305, December 2017, <<https://www.rfc-editor.org/info/rfc8305>>.
- [RFC8505] Thubert, P., Ed., Nordmark, E., Chakrabarti, S., and C. Perkins, "Registration Extensions for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Neighbor Discovery", RFC 8505, DOI 10.17487/RFC8505, November 2018, <<https://www.rfc-editor.org/info/rfc8505>>.
- [RFC8981] Gont, F., Krishnan, S., Narten, T., and R. Draves, "Temporary Address Extensions for Stateless Address Autoconfiguration in IPv6", RFC 8981, DOI 10.17487/RFC8981, February 2021, <<https://www.rfc-editor.org/info/rfc8981>>.

Acknowledgements

Thanks to the following people (in alphabetical order) for their comments, review, and feedback: Mikael Abrahamsson, Stewart Bryant, Lorenzo Colitti, Roman Danyliw, Owen DeLong, Martin Duke, Igor Gashinsky, Carles Gomez, Fernando Gont, Tatuya Jinmei, Benjamin Kaduk, Scott Kelly, Erik Kline, Warren Kumari, Barry Leiba, Jordi Palet Martinez, Erik Nordmark, Michael Richardson, Dan Romascanu, Zaheduzzaman Sarker, Michael Scharf, John Scudder, Mark Smith, Dave Thaler, Pascal Thubert, Loganaden Velvindron, and ric Vyncke.

Author's Address

Jen Linkova
Google
1 Darling Island Rd
Pyrmont NSW 2009
Australia

Email: furry@google.com