

IPv6 Maintenance
Internet-Draft
Updates: 4861 (if approved)
Intended status: Standards Track
Expires: May 29, 2020

J. Linkova
Google
November 26, 2019

Gratuitous Neighbor Discovery: Creating Neighbor Cache Entries on First-
Hop Routers
draft-linkova-6man-grand-01

Abstract

Neighbor Discovery (RFC4861) 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 [RFC4861] to allow routers to proactively create a Neighbor Cache entry when a new IPv6 address is assigned to a host. It also updates [RFC4862] and recommends hosts to send unsolicited Neighbor Advertisements upon assigning a new IPv6 address. The proposed change will minimize the delay and packet loss when a host initiate connections to off-link destination from a new IPv6 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 May 29, 2020.

Copyright Notice

Copyright (c) 2019 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	2
1.1. Requirements Language	3
1.2. Terminology	3
2. Proposed Changes to Neighbor Discovery	4
2.1. Hosts Sending Gratuitous Neighbor Advertisements	4
2.2. Routers Creating Cache Entries Upon Receiving Unsolicited Neighbor Advertisements	4
3. Avoiding Disruption	5
3.1. Neighbor Cache Entry Exists in Any State Other Than INCOMPLETE	5
3.2. Neighbor Cache Entry Does Not Exist	5
3.3. Neighbor Cache Entry is in INCOMPLETE state	6
4. Modifications to RFC-Mandated Behavior	6
4.1. Modification to RFC4861 Neighbor Discovery for IP version 6 (IPv6)	6
4.1.1. Modification to the section 7.2.5	6
4.1.2. Modification to the section 7.2.6	7
5. IANA Considerations	8
6. Security Considerations	8
7. Acknowledgements	8
8. References	9
8.1. Normative References	9
8.2. Informative References	9
Author's Address	10

1. Introduction

The Neighbor Discovery state machine defined in [RFC4861] implies that communications between IPv6 nodes are in most cases bi-directional and if a host A is trying to communicate to its neighbor, host B, the return traffic flows could be expected. So when the host A starts the address resolution process, the target host would also create an entry for the host A address in its neighbor cache. That entry will be used for sending the return traffic to the host A.

However when a host sends traffic to off-link destinations the different scenario is observed. After 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 [RFC4861] Section 7.2.2. All subsequent packets arriving before the resolution process finishes are likely to be dropped. It might cause user-visible packet loss and performance degradation

The detailed problem statement and various solution approaches could be found in [I-D.ietf-v6ops-nd-cache-init]. This document summarized the proposed neighbor discovery updates to address the issue.

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

ND: Neighbor Discovery, [RFC4861].

SLAAC: IPv6 Stateless Address Autoconfiguration, [RFC4862].

NS: Neighbor Solicitation, [RFC4861].

NA: Neighbor Advertisement, [RFC4861].

RS: Router Solicitation, [RFC4861].

RA: Router Advertisement, [RFC4861].

LLA: Link-Layer Address.

SLLA: Source link-layer Address, an option in the ND packets containing the link-layer address of the sender of the packet ([RFC4861]).

TLLA: Target link-layer Address, 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].

Optimistic DAD: a modification of DAD, [RFC4429].

2. Proposed Changes to Neighbor Discovery

The following changes are proposed to minimize the delay in creating new entries in a router neighbor cache

- o A host SHOULD send unsolicited NAs upon assigning a new IPv6 address to its interface.
- o A router SHOULD create a new cache entry upon receiving an unsolicited NA from a host.

The following sections discuss these changes in more detail.

2.1. Hosts Sending Gratuitous Neighbor Advertisements

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

To minimize the potential disruption in case of duplicate addresses the host SHOULD NOT set the Override flag for a preferred address and MUST NOT set the Override flag if the address is in Optimistic [RFC4429] state.

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 SHOULD be sent to all-routers multicast address (ff02::2). Limiting the recipients to routers only would help reduce the multicast noise level.

2.2. Routers Creating Cache Entries Upon Receiving Unsolicited Neighbor Advertisements

The 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 onlink host-to-host communication but, as

discussed in [I-D.ietf-v6ops-nd-cache-init] does not really apply for the scenario when the host is announcing its address to routers. Therefore it would be beneficial to allow routers creating new entries upon receiving an unsolicited Neighbor Advertisement.

This document suggests that routers SHOULD create a new Neighbor Cache entry when receive an unsolicited Neighbor Advertisement.

3. Avoiding Disruption

If hosts following the recommendations in this document are using the DAD mechanism defined in [RFC4862], they would send unsolicited NA as soon as the address changes the state from tentative to preferred (after its uniqueness has been verified). However hosts 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. The 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 traffic interruption of the rightful address owner in case of address conflict. As hosts willing to speed up their network stack configuration are most likely to be affected by the problem outlined in this document it seems reasonable for such hosts to advertise their optimistic GUAs 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 duplicated.

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

If the router 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 the its Optimistic address the router cache entry would not be updated with the new Link-Layer address and no impact to the traffic for the rightful address owner is expected.

3.2. Neighbor Cache Entry Does Not Exist

If there is no entry then it would be created/updated with the supplied LLA and its state set to STALE. In that case as soon as the entry is used for sending traffic to the host, the entry state will be changed to DELAY and the Neighbor Unreachability Detection would be started and the rightful owner LLA will be entered in the cache. So in the scenario when the rightful owner does not use the address for communication then it might be a short (a few seconds) period of

time when the data packets sent from the outside could reach the host with the optimistic address. However it seems likely that hosts using Optimistic DAD would start sending/receiving traffic right away, so the first return packet would trigger the NUD process and rewrite the cache.

3.3. Neighbor Cache Entry is in INCOMPLETE state

Another corner case is the INCOMPLETE cache entry for the address. If the host sends an unsolicited NA from the Optimistic address it would update the entry with the host LLA and set the entry to the STALE state. As the INCOMPLETE entry means that the router has started the ND process for the address and the multicast NS has been sent, the rightful owner is expected to reply with solicited NA with the Override flag set. Upon receiving a solicited NA with the Override flag the cache entry will be updated with the TLLA supplied and (as the NA has the Solicited flag set), the entry state will be set to REACHABLE. IT would would recover the cache entry and set the LLA to the one of the rightful owner. The only potential impact would be for packets arriving to the router after the unsolicited NA from the host but before the rightful owner responded with the solicited NA. Those packets would be sent to the host with the optimistic address instead of its rightful owner. However those packets would have been dropped anyway as until the solicited NA is received the router can not send the traffic.

4. Modifications to RFC-Mandated Behavior

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

4.1. Modification to RFC4861 Neighbor Discovery for IP version 6 (IPv6)

4.1.1. Modification to the section 7.2.5

This document proposes the following changes to the section 7.2.5 of [RFC4861]:

OLD TEXT:

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.

NEW TEXT:

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 its reachability state MUST also be set to STALE. If the received Neighbor Advertisement does not contain the Target link-layer address option the advertisement SHOULD be silently discarded.

4.1.2. Modification to the section 7.2.6

This document proposes the following changes to the section 7.2.6 of [RFC4861]:

OLD TEXT:

In such cases, a node MAY send up to MAX_NEIGHBOR_ADVERTISEMENT unsolicited Neighbor Advertisement messages to the all-nodes multicast address. These advertisements MUST be separated by at least RetransTimer seconds.

NEW TEXT:

In such cases, a node MAY send up to MAX_NEIGHBOR_ADVERTISEMENT unsolicited Neighbor Advertisement messages to the all-nodes multicast address. These advertisements MUST be separated by at least RetransTimer seconds.

A host 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 host 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:

- o if Optimistic DAD [RFC4429] is used: a new Optimistic GUA is assigned to the host interface.
- o if Optimistic DAD is not used: a GUA changes the state from tentative to preferred.

5. IANA Considerations

This memo asks the IANA for no new parameters.

6. Security Considerations

One of the potential attack vectors to consider is a cache spoofing when 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 ND cache for the given IPv6 address already exists. Therefore it is not possible for the attacker to override any existing cache 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 this document does not increase the risk of such an attack: the attacker could do it, for example, by sending a NS or RS packet with SLLAO included. All recommendations from [RFC6583] still apply.

Announcing a new address to 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 it falls into 'Security through obscurity' category. If peer-to-peer onlink communications are not desirable they should be prevented by proper layer2 security mechanisms. Therefore the risk of allowing hosts to send unsolicited Neighbor Advertisements to all-routers multicast address is low.

7. Acknowledgements

Thanks to the following people (in alphabetical order) for their comments, review and feedback: Lorenzo Colitti, Tatuya Jinmei, Erik Kline, Warren Kumari, Erik Nordmark, Michael Richardson, Dave Thaler, Pascal Thubert, Loganaden Velvindron, Eric Vyncke.

8. References

8.1. 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>>.

8.2. Informative References

- [I-D.ietf-v6ops-nd-cache-init] Linkova, J., "Neighbor Cache Entries on First-Hop Routers: Operational Considerations", draft-ietf-v6ops-nd-cache-init-00 (work in progress), October 2019.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, DOI 10.17487/RFC4941, September 2007, <<https://www.rfc-editor.org/info/rfc4941>>.
- [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>>.

Author's Address

Jen Linkova
Google
1 Darling Island Rd
Pyrmont, NSW 2009
AU

Email: furry@google.com