IPv6 Operations (v6ops) Working Group
Internet Draft
Intended status: Informational
Expires: Nov. 2025

X. Xiao
E. Vasilenko
Huawei Technologies
E. Metz
KPN
G. Mishra
Verizon Inc.
N. Buraglio
Energy Sciences Network

May 26, 2025

Neighbor Discovery Considerations in IPv6 Deployments
draft-ietf-v6ops-nd-considerations-14

Abstract

The Neighbor Discovery (ND) protocol is a critical component of the IPv6 architecture. The protocol uses multicast in many messages. It also assumes a security model where all nodes on a link are trusted. Such a design might be inefficient in some scenarios (e.g., use of multicast in wireless networks) or when nodes are not trustworthy (e.g., public access networks). These security and operational issues and the associated mitigation solutions are documented in more than 20 RFCs. There is a need to track these issues and solutions in a single document.

To that aim, this document summarizes the published ND issues and then describes how all these issues originate from three causes. Addressing the issues is made simpler by addressing the causes. This document also analyzes the mitigation solutions and demonstrates that isolating hosts into different subnets and links can help to address the three causes. Guidance is provided for selecting a suitable isolation method to prevent potential ND issues.

Status of this Memo

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 in Nov. 2025.

Table of Contents

1. Introduction

Neighbor Discovery (ND) [RFC4861] specifies the mechanisms that IPv6 nodes (hosts and routers) on the same link use to communicate and learn about each other. Stateless Address Autoconfiguration (SLAAC) [RFC4862] builds on those ND mechanisms to let nodes configure their own IPv6 addresses. When analyzing the issues nodes may encounter with ND, it helps to view the ND messages they exchange throughout their life-cycle, taking SLAAC into consideration.

For a host, the overall procedure is as follows:

1. LLA DAD: The host forms a Link-Local Address (LLA) and performs Duplicate Address Detection (DAD) using multicast Neighbor Solicitations (NSs).
2. Router Discovery: The host sends multicast Router Solicitations (RSs) to discover a router on the link. The router responds with Router Advertisements (RAs), providing subnet prefixes and other information. The host installs a Neighbor Cache Entry (NCE) for that router upon receiving the RAs. In contrast, the router cannot install an NCE for the host at this moment of the exchange because the host’s global IP address is still unknown. When the router later needs to forward a packet to the host’s global address, it will perform address resolution and install an NCE for the host.
3. GUA DAD: The host forms a Global Unicast Address (GUA) [RFC3587] or a Unique Local Address (ULA) [RFC4193] and uses

multicast NSs for DAD. For simplicity of description, this document will not further distinguish GUA and ULA.

4. Next-hop determination and address resolution: When the host needs to send a packet, it will first determine whether the next-hop is a router or an on-link host (which is the destination). If the next-hop is a router, the host already has the NCE for that router. If the next-hop is an on-link host, it will use multicast NSs to perform address resolution for the destination host. As a result, the source host installs an NCE for the destination host.
5. Node Unreachability Detection (NUD): The host uses unicast NSs to determine whether another node with an NCE is still reachable.
6. Link-layer address change announcement: If a host’s link-layer address changes, it may use multicast Node Advertisements (NAs) to announce its new link-layer address to other nodes.

For a router, the procedure is similar except that there is no Router Discovery. Instead, routers perform a Redirect procedure that hosts do not have. A router sends a Redirect to inform a node of a better next-hop for the node’s traffic.

ND uses multicast in many messages, trusts messages from all nodes, and routers may install NCEs for hosts on demand when they are to forward packets to these hosts. These may lead to issues. Concretely, various ND issues and mitigation solutions have been published in more than 20 RFCs, including:

. ND Trust Models and Threats [RFC3756],
. Secure ND [RFC3971],
. Cryptographically Generated Addresses [RFC3972],
. ND Proxy [RFC4389],
. Optimistic ND [RFC4429],
. ND for mobile broadband [RFC6459][RFC7066],
. ND for fixed broadband [TR177],
. ND Mediation [RFC6575],
. Operational ND Problems [RFC6583],
. Wireless ND (WiND) [RFC6775][RFC8505][RFC8928][RFC8929][SND],
. DAD Proxy [RFC6957],
. Source Address Validation Improvement [RFC7039],
. Router Advertisement Guard [RFC6105][RFC7113],
. Enhanced Duplicate Address Detection [RFC7527],
. Scalable ARP [RFC7586],
. Reducing Router Advertisements [RFC7772],
. Unique Prefix Per Host [RFC8273],

. ND Optimization for Transparent Interconnection of Lots of Links (TRILL) [RFC8302],
. Gratuitous Neighbor Discovery [RFC9131],
. Proxy ARP/ND for EVPN [RFC9161], and
. Using DHCPv6 Prefix Delegation (DHCPv6-PD) to Allocate Unique IPv6 Prefixes per Client in Large Broadcast Networks [RFC9663].

This document summarizes these RFCs into a one-stop reference (as of the time of writing) for easier access. This document also identifies three causes of the issues and defines three host isolation methods to address the causes and prevent potential ND issues.

1.1. Terminology

This document uses the terms defined in [RFC4861]. Additional terms are defined in this section.

MAC - To avoid confusion with link-local addresses, link-layer addresses are referred to as MAC addresses in this document.

Host Isolation - separating hosts into different subnets or links.

L3 Isolation - allocating a unique prefix per host [RFC8273][RFC9663] so that every host is in a different subnet. Given that a unique prefix can be allocated per host on shared media, hosts in different subnets may be on the same link.

L2 Isolation - taking measures to prevent a host from reaching other hosts directly in Layer 2 (L2) so that every host is in a different link. Due to the existence of Multi-Link Subnet [RFC4903], hosts in different links may be in the same subnet. Therefore, L2 Isolation does not imply L3 Isolation, and L3 Isolation does not imply L2 Isolation either.

L3+L2 Isolation - applying L3 Isolation and L2 Isolation simultaneously so that every host is in a different subnet and on a different link.

Partial L2 Isolation - using an L3 ND proxy [RFC4389] device to represent the hosts behind it to other hosts in the same subnet. Within the subnet, ND multicast exchange is segmented into multiple smaller scopes, each represented by an ND proxy device.

2. Review of Inventoried ND Issues

2.1. Multicast May Cause Performance and Reliability Issues

In some cases, ND uses multicast for NSs, NAs, RSs, and RAs. While multicast can be highly efficient in certain scenarios, e.g., in wired networks, multicast can also be inefficient in other scenarios, e.g., in large L2 networks or wireless networks.

Typically, multicast can create a large amount of protocol traffic in large L2 networks. This can consume network bandwidth, increase processing overhead, and degrade network performance [RFC7342].

In wireless networks, multicast can be inefficient or even unreliable due to a higher probability of transmission interference, lower data rate, and lack of acknowledgements (Section 3.1 of [RFC9119]).

Multicast-related performance issues of the various ND messages are summarized below:

- Issue 1: LLA DAD Degrading Performance - in an L2 network of N addresses (which can be much larger than the number of hosts, as each host can have multiple addresses), there can be N such multicast messages. This may cause performance issues when N is large.
- Issue 2: Router’s Periodic Unsolicited RAs Draining Hosts’ Battery - multicast RAs are generally limited to one packet every MIN_DELAY_BETWEEN_RAS (3 seconds), and there are usually only one or two routers on the link, so it is unlikely to cause a performance issue. However, for battery-powered hosts, such messages may wake them up and drain their batteries [RFC7772].
- Issue 3: GUA DAD Degrading Performance - same as in Issue 1.
- Issue 4: Router’s Address Resolution for Hosts Degrading Performance - same as in Issue 1.
- Issue 5: Host’s Address Resolution for Hosts Degrading Performance - same as in Issue 1.
- (For Further Study) Hosts’ MAC Address Change NAs Degrading Performance - with randomized and changing MAC addresses [MADINAS], there may be many such multicast messages.

In wireless networks, multicast is more likely to cause packet loss. Because DAD treats no response as no duplicate address detected, packet loss may cause duplicate addresses to be undetected. Multicast reliability issues are summarized below:

- Issue 6: LLA DAD Not Completely Reliable in Wireless Networks.
- Issue 7: GUA DAD Not Completely Reliable in Wireless Networks.

Note: IPv6 address collisions are extremely unlikely. As a result, these two issues are largely theoretical rather than practical.

2.2. Trusting-all-hosts May Cause On-link Security Issues

In scenarios such as public access networks, some nodes may not be trustworthy. An attacker on the link can cause the following on-link security issues [RFC3756][RFC9099]:

- Issue 8: Source IP Address Spoofing - an attacker can use another node's IP address as the source address of its ND message to pretend to be that node. The attacker can then launch various Redirect or Denial-of-Service (DoS) attacks.
- Issue 9: Denial of DAD - an attacker can repeatedly reply to a victim's DAD messages, causing the victim's address configuration procedure to fail, resulting in a DoS to the victim.
- Issue 10: Rogue RAs - an attacker can send RAs to victim hosts to pretend to be a router. The attacker can then launch various Redirect or DoS attacks.
- Issue 11: Spoofed Redirects - an attacker can send forged Redirects to victim hosts to redirect their traffic to the legitimate router itself.
- Issue 12: Replay Attacks - an attacker can capture valid ND messages and replay them later.

2.3. Router-NCE-on-Demand May Cause Forwarding Delay, NCE Exhaustion, and Address Accountability Issues

When a router needs to forward a packet to a node but does not yet have a Neighbor-Cache Entry (NCE) for that node, it first creates an NCE in the INCOMPLETE state. The router then multicasts an NS to the node's solicited-node multicast address. When the destination replies with an NA containing its MAC address, the router updates the NCE with that address and changes its state to REACHABLE, thereby completing the entry. This process is referred to as Router-NCE-on-Demand in this document.

Router-NCE-on-Demand can cause the following issues:

- Issue 13: NCE Exhaustion - an attacker can send a high volume of packets targeting non-existent IP addresses, causing the router to create numerous NCEs in the INCOMPLETE state. The

resulting resource exhaustion may cause the router to malfunction. This vulnerability, described as "NCE Exhaustion" in this document, does not require the attacker to be on-link.

. Issue 14: Router Forwarding Delay - when a packet arrives at a router, the router buffers it while attempting to determine the host's MAC address. This buffering delays forwarding and, depending on the router's buffer size, may lead to packet loss. This delay is referred to as "Router-NCE-on-Demand Forwarding Delay" in this document.

. Issue 15: Lack of Address Accountability - with SLAAC, hosts generate their IP addresses. The router does not become aware of a host's IP address until an NCE entry is created. With DHCPv6 [RFC8415], the router may not know the host's addresses unless it performs DHCPv6 snooping. In public access networks, where subscriber management often relies on IP address (or prefix) identification, this lack of address accountability poses a challenge [AddrAcc]. Without knowledge of the host's IP address, network administrators are unable to effectively manage subscribers, which is particularly problematic in public access networks. Moreover, once a router has created its NCEs, ND [RFC4861] provides no mechanism to retrieve them for management or monitoring, as noted in Section 2.6.1 of [RFC 9099].

2.4. Summary of ND Issues

The ND issues, as discussed in Sections 2.1 to 2.3, are summarized below. These issues stem from three primary causes: multicast, Trusting-all-nodes, and Router-NCE-on-Demand. Eliminating any of these causes would also mitigate the corresponding issues. These observations provide guidance for addressing and preventing ND-related issues.

(1) Multicast-related issues:

. Performance issues
  - o Issue 1: LLA DAD Degrading Performance.
  - o Issue 2: Unsolicited RA Draining Host Battery Life.
  - o Issue 3: GUA DAD degrading performance.
  - o Issue 4: Router Address Resolution for Hosts Degrading Performance.
  - o Issue 5: Host Address Resolution for Other Hosts Degrading Performance.

. Reliability issues
  - o Issue 6: LLA DAD Not Completely Reliable in Wireless Networks

- Issue 7: GUA DAD Not Completely Reliable in Wireless Networks

(2) Trusting-all-nodes related issues:

- Issue 8: Source IP Address Spoofing
- Issue 9: Denial of DAD
- Issue 10: Rogue RAs
- Issue 11: Spoofed Redirects
- Issue 12: Replay Attacks

(3) Router-NCE-on-Demand related issues:

- Issue 13: NCE Exhaustion
- Issue 14: Router Forwarding Delay
- Issue 15: Lack of Address Accountability

These issues are potential vulnerabilities and may not manifest in all usage scenarios.

When these issues may occur in a specific deployment, it is advisable to consider the mitigation solutions available. They are described in the following section.

3. Review of ND Mitigation Solutions

Table 1 summarizes ND mitigation solutions available for Issues 1-15 described in Section 2.4. Similar solutions are grouped, beginning with those that address the most issues. Unrelated solutions are ordered based on the issues (listed in Section 2.4) they address. Each solution in the table will be explained in a sub-section later, where abbreviations in the table are described.

In the table, a letter code indicates the RFC category of the mitigation solution (see BCP 9 [RFC2026] for explanation of these categories):

S - Standards Track (Proposed Standard, Draft Standard, or Internet Standard)
E - Experimental
I - Informational
B - Best Current Practice
U - Unknown (not formally defined by the IETF)

ND solu-	RFC ty-	Multicast performance	Reli- ability	On-link securi.	NCE Exhau.	Fwd. Delay	No A. Acct.

tion	pe	1	2	3	4	5	6	7	8-12	13	14	15
MBBv6	I	All identified issues solved										
FBBv6	U	All identified issues solved										
UPPH	I		X		X	X		X		X	X	X
WiND	S	All issues solved for Low-Power and Lossy Networks (LLNs)										
SARP	E					X						
ND TRILL	S					X						
ND EVPN	S					X						
7772	B		X									
GRAND	S				X						X	
SAVI/ RA G/G+	I								X			
6583	I									X		
9686	S											X

Table 1. Solutions for identified issues

3.1. ND Solution in Mobile Broadband IPv6

The IPv6 solution defined in "IPv6 in 3GPP EPS" [RFC6459], "IPv6 for 3GPP Cellular Hosts" [RFC7066], and "Extending an IPv6 /64 Prefix from a Third Generation Partnership Project (3GPP) Mobile Interface to a LAN Link" [RFC7278] is called Mobile Broadband IPv6 (MBBv6) in this document. They are Informational RFCs. The key points are:

- Putting every host, e.g., the mobile User Equipment (UE), in a Point-to-Point (P2P) link with the router, e.g., the mobile gateway. Consequently:
  - All multicast is effectively turned into unicast.
  - The P2P links do not have a MAC address. Therefore, Router-NCE-on-Demand is not needed.

- Trusting-all-nodes is only relevant to the router. By applying filtering at the router, e.g., dropping RAs from the hosts, even malicious hosts cannot cause harm.

. Assigning a unique /64 prefix to each host. Together with the P2P link, this puts each host on a separate link and subnet.
. Maintaining (prefix, interface) binding at the router for forwarding purposes.

Since all the three causes of ND issues are addressed, all the issues discussed in Section 2.4 are addressed.

3.2. ND Solution in Fixed Broadband IPv6

The IPv6 solution defined in "IPv6 in the context of TR-101" [TR177] is called Fixed Broadband IPv6 (FBBv6) in this document. FBBv6 has two flavors:

. P2P: Every host, e.g., the Residential Gateway (RG), is in a P2P link with the router, e.g., the Broadband Network Gateway (BNG). In this case, the solution is functionally similar to MBBv6. All ND issues discussed in Section 2.4 are solved.
. Point-to-Multi-Point (P2MP): All hosts, e.g., the RGs, connected to an access device, e.g., the Optical Line Terminal (OLT), are in a P2MP link with the router, e.g., the BNG. This is achieved by placing all hosts in a single VLAN on the router and configuring the OLT to block any frame from being forwarded between its access ports; traffic from each host can travel only up toward the router, not sideways to another host, thereby preventing direct host-to-host communication.

The following summarizes the two key aspects of the FBBv6-P2MP architecture as described in [TR177] and the associated benefits:

. Implementing DAD Proxy [RFC6957]:

  In a P2MP architecture described above, the normal ND DAD procedure will breaks down because hosts cannot exchange NSs with one another. To address this, the router participates in the DAD process as a DAD Proxy to resolve address duplication.

  The benefits are:

  - Multicast traffic from all hosts to the router is effectively converted into unicast, as hosts can only communicate directly with the router.

- The Trusting-all-nodes model is limited to the router. By applying simple filtering, e.g., dropping RAs from hosts, the router can mitigate security risks, even from malicious hosts

. Assigning a unique /64 prefix to each host:

Assigning each host a unique /64 prefix results in several operational improvements:

- The router can proactively install a forwarding entry for that prefix towards the host, eliminating the need for Router-NCE-on-Demand.
- Since each host resides in a different subnet, traffic between hosts is routed through the router, eliminating the need for hosts to perform address resolution for one another.
- Without address resolution, router multicast to hosts is limited to unsolicited RAs. As each host resides in its own subnet, these RAs are sent as unicast packets to individual hosts. This follows the approach specified in [RFC6085], where the host’s MAC address replaces the multicast MAC address in the RA.

Since all three causes of ND issues are addressed, all ND issues (Section 2.4) are also addressed.

3.3. Unique IPv6 Prefix per Host (UPPH)

UPPH solutions are described in [RFC8273] and [RFC9663]. Both are Informational RFCs. [RFC8273] relies on SLAAC for unique prefix allocation while [RFC9663] relies on DHCP-PD. That difference in allocation mechanism does not change the discussion on ND issues, because every IPv6 node is still required to run SLAAC, even when it receives its prefix via DHCP-PD. Therefore, discussing [RFC8273] alone is sufficient.

[RFC8273] "improves host isolation and enhanced subscriber management on shared network segments" such as Wi-Fi or Ethernet. The key points are:

. When a prefix is allocated to the host, the router can proactively install a forwarding entry for that prefix towards the host. There is no more Router-NCE-on-Demand.

- Without address resolution, router multicast to hosts consists only of unsolicited RAs. They will be sent to hosts one by one in unicast because the prefix for every host is different.
- Since different hosts are in different subnets, hosts will send traffic to other hosts via the router. There is no host-to-host address resolution.

Therefore, ND issues caused by Router-NCE-on-Demand and router multicast to hosts are prevented.

[RFC8273] indicates that a "network implementing a unique IPv6 prefix per host can simply ensure that devices cannot send packets to each other except through the first-hop router". But when hosts are on a shared medium like Ethernet, ensuring "devices cannot send packets to each other except through the first-hop router" requires additional measures like Private VLAN [RFC5517]. Without such additional measures, on a shared medium, hosts can still reach each other in L2 as they belong to the same Solicited-Node Multicast Group. Therefore, Trusting-all-nodes and host multicast to routers may cause issues. Of the host multicast issues (i.e., Issues 1, 3, 5, 6, and 7), Unique Prefix per Host prevents Issues 5 and 7, because there is no need for address resolution among hosts (Issue 5) and there is no possibility of GUA duplication (Issue 7). But Issues 1, 3, and 6 may occur.

3.4. Wireless ND and Subnet ND

Wireless ND (WiND) [RFC6775][RFC8505][RFC8928][RFC8929] (Standards Track) defines a fundamentally different ND solution for Low-Power and Lossy Networks (LLNs) [RFC7102]. WiND changes host and router behaviors to use multicast only for router discovery. The key points are:

- Hosts use unicast to proactively register their addresses at the routers. Routers use unicast to communicate with hosts and become an abstract registrar and arbitrator for address ownership.
- The router also proactively installs NCEs for the hosts. This avoids the need for address resolution for the hosts.
- The router sets Prefix Information Option (PIO) L-bit to 0. Each host communicates only with the router (Section 6.3.4 of [RFC4861]).
- Other functionalities that are relevant only to LLNs.

WiND addresses all ND issues (Section 2.4) in LLNs. However, WiND support is not mandatory for general-purpose hosts. Therefore, it

cannot be relied upon as a deployment option without imposing additional constraints on the participating nodes.

3.5. Scalable Address Resolution Protocol

Scalable Address Resolution Protocol [RFC7586] was an Experimental solution. That experiment ended in 2017, two years after the RFC was published. Because the idea has been used in mitigation solutions for more specific scenarios (described in Sections 3.6 and 3.7), it is worth describing here. The usage scenario is Data Centers (DCs), where large L2 domains span across multiple sites. In each site, multiple hosts are connected to a switch. The hosts can be Virtual Machines (VMs), so the number can be large. The switches are interconnected by a native or overlay L2 network.

The switch will snoop and install (IP, MAC address) proxy table for the local hosts. The switch will also reply to address resolution requests from other sites to its hosts with its own MAC address. In doing so, all hosts within a site will appear to have a single MAC address to other sites. As such, a switch only needs to build a MAC address table for the local hosts and the remote switches, not for all the hosts in the L2 domain. Consequently, the MAC address table size of the switches is significantly reduced. A switch will also add the (IP, MAC address) replies from remote switches to its proxy ND table so that it can reply to future address resolution requests from local hosts for such IPs directly. This greatly reduces the number of address resolution multicast in the network.

Unlike MBBv6, FBBv6, and UPPH, which try to address all ND issues discussed in Section 2.4, SARP focuses on reducing address resolution multicast to improve the performance and scalability of large L2 domains in DCs.

3.6. ARP and ND Optimization for TRILL

ARP and ND Optimization for TRILL [RFC8302] (Standards Track) is similar to SARP (Section 3.5). It can be considered an application of SARP in the TRILL environment.

Like SARP, ARP, and ND Optimization for TRILL focuses on reducing multicast address resolution. That is, it addresses Issue 5 (Section 2.1).

3.7. Proxy ARP/ND in Ethernet Virtual Private Networks (EVPN)

Proxy ARP/ND in EVPN is specified in [RFC9161] (Standards Track). The usage scenario is DCs where large L2 domains span across multiple sites. In each site, multiple hosts are connected to a Provider Edge (PE) router. The PEs are interconnected by EVPN tunnels.

PE of each site snoops the local address resolution NAs to build (IP, MAC address) Proxy ND table entries. PEs then propagate such Proxy ND entries to other PEs via the Border Gateway Protocol (BGP). Each PE also snoops local hosts' address resolution NSs for remote hosts. If an entry exists in its Proxy ND table for the remote hosts, the PE will reply directly. Consequently, the number of multicast address resolution messages is significantly reduced.

Like SARP, Proxy ARP/ND in EVPN also focuses on reducing address resolution multicast.

3.8. Reducing Router Advertisements

Maintaining IPv6 connectivity requires that hosts be able to receive periodic multicast RAs [RFC4861]. Hosts that process unicast packets while they are asleep must also process multicast RAs while they are asleep. An excessive number of RAs can significantly reduce the battery life of mobile hosts. [RFC7772] (Best Current Practice) specifies a solution to reduce RAs:

- The router should respond to RS with unicast RA (rather than the normal multicast RA) if the host's source IP address is specified and the host's MAC address is valid. This way, other hosts will not receive this RA.
- The router should reduce the multicast RA frequency

[RFC7772] addresses Issue 2 (Section 2.1).

3.9. Gratuitous Neighbor Discovery (GRAND)

GRAND [RFC9131] (Standards Track) changes ND in the following ways:

- A node sends unsolicited NAs upon assigning a new IPv6 address to its interface.
- A router creates a new NCE for the node and sets its state to STALE.

When a packet for the host later arrives, the router can use the existing STALE NCE to forward it immediately ([RFC4861] Section 7.2.2). It then verifies reachability by sending a unicast NS rather than a multicast one for address resolution. In this way, GRAND eliminates the Router Forwarding Delay. But it does not solve other Router-NCE-on-Demand issues. For example, NCE Exhaustion can still happen.

3.10. Source Address Validation Improvement (SAVI) and Router Advertisement Guard

SAVI [RFC7039] (Informational) binds an address to a port on an L2 switch and rejects claims from other ports for that address. Therefore, a node cannot spoof the IP address of another node.

Router Advertisement Guard (RA-Guard) [RFC6105][RFC7113] (Informational) only allows RAs from a port that a router is connected to. Therefore, nodes on other ports cannot pretend to be a router.

SAVI and RA-Guard address the on-link security issues.

3.11. RFC 6583 Dealing with NCE Exhaustion Attacks

[RFC6583] (Informational) deals with the NCE Exhaustion attack issue (Section 2.3). It recommends that:

- Operators should
  - Filter unused address space so that messages to such addresses can be dropped rather than triggering NCE creation.
  - Implement rate-limiting mechanisms for ND message processing to prevent CPU and memory resources from being overwhelmed.
- Vendors should
  - Prioritizing NDP processing for existing NCEs over creating new NCEs

[RFC6583] acknowledges that "some of these options are ’kludges’, and can be operationally difficult to manage". [RFC6583] partially addresses the Router NCE Exhaustion issue. In practice, router vendors cap the number of NCEs per interface to prevent cache exhaustion. If the link has more addresses than that cap, the router cannot keep an entry for every address, and packets destined for addresses without an NCE are simply dropped [RFC9663].

3.12. Registering Self-generated IPv6 Addresses using DHCPv6

In IPv4, network administrators can retrieve a host's IP address from the DHCP server and use it for subscriber management. In IPv6 and SLAAC, this is not possible (Section 2.3).

[RFC9686] (Standards Track) defines a method for informing a DHCPv6 server that a host has one or more self-generated or statically configured addresses. This enables network administrators to retrieve the IPv6 addresses for each host from the DHCPv6 server. [RFC9686] provides a solution for Issue 15 (Section 2.3).

3.13. Enhanced DAD

Enhanced DAD [RFC7527] (Standards Track) addresses a DAD failure issue in a specific situation: a looped back interface. DAD will fail in a looped-back interface because the sending host will receive the DAD message back and will interpret it as another host is trying to use the same address. The solution is to include a Nonce option [RFC3971] in each DAD message so that the sending host can detect that the looped-back DAD message is sent by itself.

Enhanced DAD does not solve any ND issue. It extends ND to work in a new scenario: looped-back interface. It is reviewed here only for completeness.

3.14. ND Mediation for IP Interworking of Layer 2 VPNs

ND mediation is specified in [RFC6575] (Standards Track). When two Attachment Circuits (ACs) are interconnected by a Virtual Private Wired Service (VPWS), and the two ACs are of different media (e.g., one is Ethernet while the other is Frame Relay), the two PEs must interwork to provide mediation service so that a Customer Edge (CE) can resolve the MAC address of the remote end. [RFC6575] specifies such a solution.

ND Mediation does not address any ND issue. It extends ND to work in a new scenario: two ACs of different media interconnected by a VPWS. It is reviewed here only for completeness.

3.15. ND Solutions Defined before the Latest Versions of ND

The latest versions of ND and SLAAC are specified in [RFC4861] and [RFC4862]. Several ND mitigation solutions were published before [RFC4861]. They are reviewed in this section only for completeness.

3.15.1. Secure Neighbor Discovery (SeND)

The purpose of SeND [RFC3971] (Standards Track) is to ensure that hosts and routers are trustworthy. SeND defined three new ND options, i.e., Cryptographically Generated Addresses (CGA) [RFC3972] (Standards Track), RSA public-key cryptosystem, and Timestamp/Nonce, an authorization delegation discovery process, an address ownership proof mechanism, and requirements for the use of these components in the ND protocol.

3.15.2. Cryptographically Generated Addresses (CGA)

The purpose of CGA is to associate a cryptographic public key with an IPv6 address in the SeND protocol. The key point is to generate the Interface Identifier (IID) of an IPv6 address by computing a cryptographic hash of the public key. The resulting IPv6 address is called a CGA. The corresponding private key can then be used to sign messages sent from the address.

CGA assumes that a legitimate host does not care about the bit combination of the IID that would be created by some hash procedure. The attacker needs an exact IID to impersonate the legitimate hosts, but then the attacker is challenged to do a reverse hash calculation which is a strong mathematical challenge.

CGA is part of SeND. There is no reported deployment.

3.15.3. ND Proxy

ND Proxy [RFC4389] (Experimental) aims to enable multiple links joined by an ND Proxy device to work as a single link.

- When an ND Proxy receives an ND request from a host on a link, it will proxy the message out the "best" (defined in the next paragraph) outgoing interface. If there is no best interface, the ND Proxy will proxy the message to all other links. Here, proxy means acting as if the ND message originates from the ND Proxy itself. That is, the ND Proxy will change the ND message’s source IP and source MAC address to the ND Proxy’s outgoing interface’s IP and MAC address, and create an NCE entry at the outgoing interface accordingly.
- When ND Proxy receives an ND reply, it will act as if the ND message is destined for itself, and update the NCE entry state at the receiving interface. Based on such state information, the ND Proxy can determine the "best" outgoing interface for

future ND requests. The ND Proxy then proxies the ND message back to the requesting host.

ND Proxy is widely used in SARP (Sections 3.5), ND Optimization for TRILL (Sections 3.6), and Proxy ARP/ND in EVPN (Sections 3.7).

3.15.4. Optimistic DAD

Optimistic DAD [RFC4429] (Standards Track) seeks to minimize address configuration delays in the successful case and to reduce disruption as far as possible in the failure case. That is, Optimistic DAD lets hosts immediately use the newly formed address to communicate before DAD completes, assuming that DAD will succeed anyway. If the address turns out to be duplicate, Optimistic DAD provides a set of mechanisms to minimize the impact. Optimistic DAD modified the original ND [RFC2461] and SLAAC [RFC2462], but the solution was not incorporated into the latest specifications of [RFC4861] and [RFC4862]. However, implementations of Optimistic DAD exist.

Optimistic DAD does not solve any ND issue (Section 2). It is reviewed here only for completeness.

4. Guidelines for Prevention of Potential ND Issues

By knowing the potential ND issues and associated mitigation solutions, network administrators of existing IPv6 deployments can assess whether these issues may occur in their networks and, if so, whether to deploy the mitigation solutions proactively. Deploying these solutions may take time and additional resources. Therefore, it is advisable to plan.

Network administrators planning to start their IPv6 deployments can use the issue-solution information to help plan their deployments. Moreover, they can take proactive action to prevent potential ND issues.

4.1. Learning Host Isolation from the Existing Solutions

While various ND solutions may initially appear unrelated, categorizing them into four distinct groups highlights an important observation: "host isolation" is an effective strategy for mitigating ND-related issues.

Group 1: L3 and L2 Isolation

This group includes MBBv6 and FBBv6, which isolate hosts at both L3 and L2 by placing each host within its subnet and link. This prevents ND issues caused by multicast and Trusting-all-nodes, as each host operates within its isolated domain. Furthermore, since routers can route packets to a host based on its unique prefix, the need for Router-NCE-on-Demand is also eliminated. Therefore, L3 and L2 Isolation prevent all ND issues.

Group 2: L3 Isolation

This group includes UPPH solutions like [RFC8273] and [RFC9663], which isolate hosts into separate subnets while potentially leaving them on the same shared medium. This approach mitigates ND issues caused by router multicast to hosts and eliminates the need for "Router-NCE-on-Demand", as detailed in Section 3.3.

Group 3: Partial L2 Isolation

This group encompasses solutions such as WiND, SARP, ND Optimization for TRILL, and Proxy ND in EVPN. These solutions use a proxy device to represent the hosts behind it, effectively isolating those hosts into distinct multicast domains. While hosts are still located within the same subnet, their separation into different multicast domains reduces the scope of ND issues related to multicast-based address resolution.

Group 4: Non-Isolating Solutions

The final group includes remaining solutions that do not implement host isolation. These solutions do not prevent ND issues but instead focus on addressing specific ND problems.

The analysis demonstrates that the stronger the isolation of hosts, the more ND issues can be mitigated. This correlation is intuitive, as isolating hosts reduces the multicast scope, minimizes the number of nodes that must be trusted, and may eliminate the need for "Router-NCE-on-Demand", the three primary causes of ND issues.

This understanding can be used to prevent ND issues.

4.2. Applicability of Various Isolation Methods

4.2.1. Applicability of L3+L2 Isolation

Benefits:

- All ND issues (Section 2.4) can be effectively mitigated.

Constraints:

1. L2 Isolation:

Actions must be taken to isolate hosts in L2. The required effort varies by the chosen method and deployment context. For example, the P2P method [RFC7066] is heavy-weight, while the Private VLAN method [RFC5517] is more manageable.

2. Unique Prefix Allocation:

A large number of prefixes will be required, with one prefix assigned per host. This is generally not a limitation for IPv6. For instance, members of a Regional Internet Registry (RIR) can obtain a /29 prefix allocation [RIPE738], which provides 32 billion /64 prefixes - sufficient for any foreseeable deployment scenarios. Practical implementations, such as MBBv6 assigning /64 prefixes to billions of mobile UEs [RFC6459] and FBBv6 assigning /56 prefixes to hundreds of millions of routed RGs [TR177], demonstrate the feasibility of this approach.

3. Privacy Issue from Unique Prefix Identifiability:

Assigning a unique prefix to each host may theoretically reduce privacy, as hosts can be directly identified by their assigned prefix. However, alternative host identification methods, such as cookies, are commonly used. Therefore, unique prefix identifiability may not make much difference. The actual impact on privacy is therefore likely to be limited.

4. Router Support for L3 Isolation:

The router must support an L3 Isolation solution, e.g., [RFC8273] or [RFC9663].

5. A Large Number of Router Interfaces May be Needed:

If the P2P method is used, the router must instantiate a separate logical interface for every attached host. In this case, a large number of interfaces will be needed at the router.

6. Router as a Bottleneck:

Since all communication between hosts is routed through the router, the router may become a performance bottleneck in high-traffic scenarios.

7. Incompatibility with Host-Based Multicast Services:

Services that rely on multicast communication among hosts, such as Multicast Domain Name System [RFC6762], will be disrupted.

4.2.2. Applicability of L3 Isolation

Benefits:

. All ND issues (Section 2.4) are mitigated, with the exception of:
  - o LLA DAD multicast degrading performance,
  - o LLA DAD not reliable in wireless networks, and
  - o On-link security

  These remaining issues depend on the characteristics of the shared medium:

  - o If the shared medium is Ethernet, the issues related to LLA DAD multicast are negligible.
  - o If nodes can be trusted, such as in private networks, on-link security concerns are not significant.

. No need for L2 Isolation. Consequently, this method can be applied in a wide range of scenarios, making it possibly the most practical host isolation method.

Constraints, as discussed in Section 4.2.1:

1. Unique Prefix Allocation

2. Router Support for L3 Isolation

3. Router as a Bottleneck

4. Privacy Issue from Unique Prefix Identifiability.

4.2.3. Applicability of Partial L2 Isolation

Benefits:

. Reduced Multicast Traffic:

This method reduces multicast traffic, particularly for address resolution, by dividing the subnet into multiple multicast domains.

Constraint:

. Router Support for Partial L2 Isolation:

The router must implement a Partial L2 Isolation solution such as WiND, SARP, ND Optimization for TRILL, and Proxy ND in EVPN to support this method.

4.3. Guidelines for Applying Isolation Methods

Based on the applicability analysis provided in the preceding sections, network administrators can determine whether to implement an isolation method and, if so, which method is most appropriate for their specific deployment.

A simple guideline is to consider the isolation methods in the order listed in the preceding sections, progressing from the strongest isolation to the weakest:

. Stronger isolation methods can prevent more ND issues, but may also impose higher entry requirements.
. Weaker isolation methods have fewer entry requirements but may leave some ND issues unmitigated.

The choice between L3+L2 Isolation and L3 Isolation often depends on the cost of implementing L2 Isolation:

. If the cost is acceptable, L3+L2 Isolation is preferable because it eliminates every ND issue listed in Section 2.4.
. Otherwise, L3 Isolation addresses most of those issues while keeping the implementation effort reasonable.

Selecting an isolation method that is either too strong or too weak does not result in serious consequences:

. Choosing an overly strong isolation method may require the network administrator to meet higher entry requirements initially, such as measures for L2 Isolation, additional prefixes, or additional router capabilities.
. Choosing a "weaker isolation method" may necessitate deploying supplemental ND mitigation techniques to address any unresolved ND issues.

In either case, the resulting solution can be functional and effective.

5. Security Considerations

This document is a review of known ND issues and solutions, including security. It does not introduce any new solutions. Therefore, it does not introduce new security issues.

6. IANA Considerations

This document has no request to IANA.

7. References

7.1. Normative References

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861.

[RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862.

7.2. Informative References

[AddrAcc] T. Chown, C. Cummings, D.Carder, "IPv6 Address Accountability Considerations", Internet draft, Oct. 2024.

[MADINAS] J. Henry, Y. Lee, "Randomized and Changing MAC Address: Context, Network Impacts, and Use Cases", draft-ietf-madinas-use-cases-19.

[RFC2026] S. Bradner, "The Internet Standards Process -- Revision 3", RFC 2026.

[RFC2461] T. Narten, E. Nordmark, W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, obsoleted by RFC 4861.

[RFC2462] S. Thomson, T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, obsoleted by RFC 4862.

[RFC3587] R. Hinden, S. Deering, E. Nordmark, "IPv6 Global Unicast Address Format", RFC 3587.

[RFC3756] P. Nikander, J. Kempf, E. Nordmark, "IPv6 Neighbor Discovery (ND) Trust Models and Threats", RFC 3756.

[RFC3971] J. Arkko, J. Kempf, B. Zill, P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC3971.

[RFC3972] T. Aura, "Cryptographically Generated Addresses (CGA)", RFC3972.

[RFC4193] R. Hinden, B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193.

[RFC4389] D. Thaler, M. Talwar, C. Patel, "Neighbor Discovery Proxies (ND Proxy)", RFC 4389.

[RFC4429] N. Moore, "Optimistic Duplicate Address Detection (DAD) for IPv6", RFC 4429.

[RFC4903] D. Thaler, "Multi-Link Subnet Issues", RFC 4903.

[RFC5517] S. HomChaudhuri, M. Foschiano, "Cisco Systems' Private VLANs: Scalable Security in a Multi-Client Environment", RFC 5517.

[RFC6085] S. Gundavelli, M. Townsley, O. Troan, W. Dec, "Address Mapping of IPv6 Multicast Packets on Ethernet", RFC 6085.

[RFC6105] E. Levy-Abegnoli, G. Van de Velde, C. Popoviciu, J. Mohacsi, "IPv6 Router Advertisement Guard", RFC 6105.

[RFC6459] J. Korhonen, J. Soininen, B. Patil, T. Savolainen, G. Bajko, K. Iisakkila, "IPv6 in 3rd Generation Partnership Project (3GPP) Evolved Packet System (EPS)", RFC 6459.

[RFC6575] H. Shah, E. Rosen, G. Heron, V. Kompella, "Address Resolution Protocol (ARP) Mediation for IP Interworking of Layer 2 VPNs", RFC 6575.

[RFC6583] I. Gashinsky, J. Jaeggli, W. Kumari, "Operational Neighbor Discovery Problems", RFC 6583.

[RFC6762] S. Cheshire, M. Krochmal, "Multicast DNS", RFC 6762.

[RFC6775] Z. Shelby, S. Chakrabarti, E. Nordmark, C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775.

[RFC6957] F. Costa, J-M. Combes, X. Pougnard, H. Li, "Duplicate Address Detection Proxy", RFC 6957

[RFC7039] J. Wu, J. Bi, M. Bagnulo, F. Baker, C. Vogt, "Source Address Validation Improvement (SAVI) Framework", RFC 7039.

[RFC7066] J. Korhonen, J. Arkko, T. Savolainen, S. Krishnan, "IPv6 for Third Generation Partnership Project (3GPP) Cellular Hosts", RFC 7066.

[RFC7102] JP. Vasseur, "Terms Used in Routing for Low-Power and Lossy Networks", RFC 7102.

[RFC7113] F. Gont, "Implementation Advice for IPv6 Router Advertisement Guard (RA-Guard)", RFC 7113.

[RFC7278] Extending an IPv6 /64 Prefix from a Third Generation Partnership Project (3GPP) Mobile Interface to a LAN Link", RFC7278.

[RFC7342] L. Dunbar, W. Kumari, I. Gashinsky, "Practices for Scaling ARP and Neighbor Discovery (ND) in Large Data Centers", RFC 7342.

[RFC7527] R. Asati, H. Singh, W. Beebee, C. Pignataro, E. Dart, W. George, "Enhanced Duplicate Address Detection", RFC 7527.

[RFC7586] Y. Nachum, L. Dunbar, I. Yerushalmi, T. Mizrahi, "The Scalable Address Resolution Protocol (SARP) for Large Data Centers", RFC7586.

[RFC7772] A. Yourtchenko, L. Colitti, "Reducing Energy Consumption of Router Advertisements", RFC 7772.

[RFC8273] J. Brzozowski, G. Van de Velde, "Unique IPv6 Prefix per Host", RFC 8273.

[RFC8302] Y. Li, D. Eastlake 3rd, L. Dunbar, R. Perlman, M. Umair, "Transparent Interconnection of Lots of Links (TRILL): ARP and Neighbor Discovery (ND) Optimization", RFC 8302.

[RFC8415] T. Mrugalski, M. Siodelski, A. Yourtchenko, M. Richardson, S. Jiang, T. Lemon, T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 8415.

[RFC8505] P. Thubert, E. Nordmark, S. Chakrabarti, C. Perkins, "Registration Extensions for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Neighbor Discovery", RFC 8505.

[RFC8928] P. Thubert, B. Sarikaya, M. Sethi, R. Struik, "Address-Protected Neighbor Discovery for Low-Power and Lossy Networks", RFC 8928.

[RFC8929] P. Thubert, C.E. Perkins, E. Levy-Abegnoli, "IPv6 Backbone Router", RFC 8929.

[RFC9099] E. Vyncke, K. Chittimaneni, M. Kaeo, E. Rey, "Operational Security Considerations for IPv6 Networks", RFC 9099.

[RFC9119] C. Perkins, M. McBride, D. Stanley, W. Kumari, JC. Zuniga, "Multicast Considerations over IEEE 802 Wireless Media", RFC 9119.

[RFC9131] J. Linkova, "Gratuitous Neighbor Discovery: Creating Neighbor Cache Entries on First-Hop Routers", RFC 9131.

[RFC9161] J. Rabadan, S. Sathappan, K. Nagaraj, G. Hankins, T. King, "Operational Aspects of Proxy ARP/ND in Ethernet Virtual Private Networks", RFC 9161.

[RFC9663] L. Colitti, J. Linkova, X. Ma, "Using DHCP-PD to Allocate Unique IPv6 Prefix per Client in Large Broadcast Networks", RFC 9663.

[RFC9686] W. Kumari, S. Krishnan, R. Asati, L. Colitti, J. Linkova, S. Jiang, "Registering Self-generated IPv6 Addresses using DHCPv6", RFC 9686.

[RIPE738] IPv6 Address Allocation and Assignment Policy, https://www.ripe.net/publications/docs/ripe-738

[SND] P. Thubert, M. Richardson, "Architecture and Framework for IPv6 over Non-Broadcast Access", Internet draft, June 2023.

[TR177] S. Ooghe, B. Varga, W. Dec, D. Allan, "IPv6 in the context of TR-101", Broadband Forum, TR-177.

8. Acknowledgments

The authors would like to thank Eric Vyncke, Gunter Van de Velde, Lorenzo Colitti, Erik Kline, Warren Kumari, Mohamed Boucadair, Gorry Fairhurst, Pascal Thubert, Jen Linkova, Brian Carpenter, Mike Ackermann, Nalini Elkins, Ed Horley, Ole Troan, David Thaler, Chongfeng Xie, Chris Cummings, Dale Carder, Tim Chown, Priyanka Sinha, Aijun Wang, Ines Robles, Magnus Westerlund, Barry Leiba, Deb Cooley and Paul Wouters for their reviews and comments. The authors would also like to thank Tim Winters for being the document shepherd.

Authors' Addresses

XiPeng Xiao
Huawei Technologies Dusseldorf
Hansaallee 205, 40549 Dusseldorf, Germany
Email: xipengxiao@huawei.com

Eduard Vasilenko
Huawei Technologies
17/4 Krylatskaya st, Moscow, Russia 121614
Email: vasilenko.eduard@huawei.com

Eduard Metz
KPN N.V.
Email: eduard.metz@kpn.com

Gyan Mishra
Verizon Inc.
Email: gyan.s.mishra@verizon.com

Nick Buraglio
Energy Sciences Network
Email: buraglio@es.net