ï»¿

Internet Engineering Task Force (IETF)
Request for Comments: 8929
Updates: 6775, 8505
Category: Standards Track
ISSN: 2070-1721

P. Thubert, Ed.
Cisco Systems
C.E. Perkins
Blue Meadow Networking
E. Levy-Abegnoli
Cisco Systems

November 2020

IPv6 Backbone Router

Abstract

This document updates RFCs 6775 and 8505 in order to enable proxy services for IPv6 Neighbor Discovery by Routing Registrars called "Backbone Routers". Backbone Routers are placed along the wireless edge of a backbone and federate multiple wireless links to form a single Multi-Link Subnet (MLSN).

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/rfc8929.

Table of Contents

1. Introduction

Ethernet bridging per IEEE Std 802.1 [IEEEstd8021Q] provides an efficient and reliable broadcast service for wired networks; applications and protocols have been built that heavily depend on that feature for their core operation. Unfortunately, Low-Power and Lossy Networks (LLNs) and local wireless networks generally do not provide the broadcast capabilities of Ethernet bridging in an economical fashion.

As a result, protocols designed for bridged networks that rely on multicast and broadcast often exhibit disappointing behaviors when employed unmodified on a local wireless medium (see [MCAST-PROBLEMS]).

Wi-Fi [IEEEstd80211] Access Points (APs) deployed in an Extended Service Set (ESS) act as Ethernet bridges [IEEEstd8021Q], with the property that the bridging state is established at the time of association. This ensures connectivity to the end node (the Wi-Fi Station (STA)) and protects the wireless medium against broadcast-intensive transparent bridging [IEEEstd8021Q] reactive lookups. In other words, the association process is used to register the link-layer address of the STA to the AP. The AP subsequently proxies the bridging operation and does not need to forward the broadcast lookups over the radio.

In the same way as transparent bridging, the IPv6 [RFC8200] Neighbor Discovery (IPv6 ND) protocol [RFC4861] [RFC4862] is a reactive protocol, based on multicast transmissions to locate an on-link correspondent and ensure the uniqueness of an IPv6 address. The mechanism for Duplicate Address Detection (DAD) [RFC4862] was designed for the efficient broadcast operation of Ethernet bridging. Since broadcast can be unreliable over wireless media, DAD often fails to discover duplications [DAD-ISSUES]. In practice, the fact that IPv6 addresses very rarely conflict is mostly attributable to the entropy of the 64-bit Interface IDs as opposed to the successful operation of the IPv6 ND DAD and resolution mechanisms.

The IPv6 ND Neighbor Solicitation (NS) [RFC4861] message is used for DAD and address lookup when a node moves or wakes up and reconnects to the wireless network. The NS message is targeted to a Solicited-Node Multicast Address (SNMA) [RFC4291] and should, in theory, only reach a very small group of nodes. But, in reality, IPv6 multicast messages are typically broadcast on the wireless medium, so they are processed by most of the wireless nodes over the subnet (e.g., the ESS fabric) regardless of how few of the nodes are subscribed to the SNMA. As a result, IPv6 ND address lookups and DADs over a large wireless network and/or LLN can consume enough bandwidth to cause a substantial degradation to the unicast traffic service.

Because IPv6 ND messages sent to the SNMA group are broadcast at the radio link layer, wireless nodes that do not belong to the SNMA group still have to keep their radio turned on to listen to multicast NS messages, which is a waste of energy for them. In order to reduce their power consumption, certain battery-operated devices such as

Internet of Things (IoT) sensors and smartphones ignore some of the broadcasts, making IPv6 ND operations even less reliable.

These problems can be alleviated by reducing the IPv6 ND broadcasts over wireless access links. This has been done by splitting the broadcast domains and routing between subnets. At the extreme, this can be done by assigning a /64 prefix to each wireless node (see [RFC8273]). But deploying a single large subnet can still be attractive to avoid renumbering in situations that involve large numbers of devices and mobility within a bounded area.

A way to reduce the propagation of IPv6 ND broadcast in the wireless domain while preserving a large single subnet is to form a Multi-Link Subnet (MLSN). Each link in the MLSN, including the backbone, is its own broadcast domain. A key property of MLSNs is that link-local unicast traffic, link-scope multicast, and traffic with a hop limit of 1 will not transit to nodes in the same subnet on a different link, which is something that may produce unexpected behavior in software that expects a subnet to be entirely contained within a single link.

This specification considers a special type of MLSN with a central backbone that federates edge (LLN) links, with each link providing its own protection against rogue access and tempering or replaying packets. In particular, the use of classical IPv6 ND on the backbone requires that the all nodes are trusted and that rogue access to the backbone is prevented at all times (see Section 11).

In that particular topology, ND proxies can be placed at the boundary of the edge links and the backbone to handle IPv6 ND on behalf of Registered Nodes and to forward IPv6 packets back and forth. The ND proxy enables the continuity of IPv6 ND operations beyond the backbone and enables communication using Global or Unique Local Addresses between any pair of nodes in the MLSN.

The 6LoWPAN Backbone Router (6BBR) is a Routing Registrar [RFC8505] that provides ND proxy services. A 6BBR acting as a Bridging Proxy provides an ND proxy function with Layer 2 continuity and can be collocated with a Wi-Fi AP as prescribed by IEEE Std 802.11 [IEEEstd80211]. A 6BBR acting as a Routing Proxy is applicable to any type of LLN, including LLNs that cannot be bridged onto the backbone, such as IEEE Std 802.15.4 [IEEEstd802154].

Knowledge of which address to proxy can be obtained by snooping the IPv6 ND protocol (see [SAVI-WLAN]), but it has been found to be unreliable. An IPv6 address may not be discovered immediately due to a packet loss or if a "silent" node is not currently using one of its addresses. A change of state (e.g., due to movement) may be missed or misordered, leading to unreliable connectivity and incomplete knowledge of the state of the network.

With this specification, the address to be proxied is signaled explicitly through a registration process. A 6LoWPAN Node (6LN) registers all of its IPv6 addresses using NS messages with an Extended Address Registration Option (EARO) as specified in [RFC8505] to a 6LoWPAN Router (6LR) to which it is directly attached. If the 6LR is a 6BBR, then the 6LN is both the Registered Node and the Registering Node. If not, then the 6LoWPAN Border Router (6LBR) that serves the LLN proxies the registration to the 6BBR. In that case, the 6LN is the Registered Node and the 6LBR is the Registering Node. The 6BBR performs IPv6 ND operations on its backbone interface on behalf of the 6LNs that have Registered Addresses on its LLN interfaces, without the need of a broadcast over the wireless medium.

A Registering Node that resides on the backbone does not register to the SNMA groups associated to its Registered Addresses and defers to the 6BBR to answer or preferably forward the corresponding multicast packets to it as unicast.

2. Terminology

2.1. 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.

2.2. New Terms

This document introduces the following terminology:

Federated: A subnet that comprises a backbone, and one or more (wireless) access links, is said to be federated into one MLSN. The ND proxy operation of 6BBRs over the backbone extends IPv6 ND operation over the access links.

Sleep Proxy: A 6BBR acts as a Sleep Proxy if it answers IPv6 ND NSs over the backbone on behalf of the Registering Node that is in a sleep state and that cannot answer in due time.

Routing Proxy: A Routing Proxy provides IPv6 ND proxy functions and enables the MLSN operation over federated links that may not be compatible for bridging. The Routing Proxy advertises its own link-layer address as the Target Link-Layer Address (TLLA) in the proxied Neighbor Advertisements (NAs) over the backbone and routes at the network layer between the federated links.

Bridging Proxy: A Bridging Proxy provides IPv6 ND proxy functions while preserving forwarding continuity at the link layer. In that case, the link-layer address and the mobility of the Registering Node is visible across the bridged backbone. The Bridging Proxy advertises the link-layer address of the Registering Node in the TLLAO in the proxied NAs over the backbone, and it proxies ND for all unicast addresses including link-local addresses. Instead of replying on behalf of the Registering Node, a Bridging Proxy will preferably forward the NS(Lookup) and Neighbor Unreachability Detection (NUD) messages that target the Registered Address to the Registering Node as unicast frames, so it can respond in its own.

Binding Table: The Binding Table is an abstract database that is maintained by the 6BBR to store the state associated with its registrations.

Binding: A Binding is an abstract state associated to one registration; in other words, it's associated to one entry in the Binding Table.

2.3. Abbreviations

This document uses the following abbreviations:

6BBR:	6LoWPAN Backbone Router
6LBR:	6LoWPAN Border Router
6LN:	6LoWPAN Node
6LR:	6LoWPAN Router
AP:	Access Point
ARO:	Address Registration Option
DAC:	Duplicate Address Confirmation
DAD:	Duplicate Address Detection
DAR:	Duplicate Address Request
DODAG:	Destination-Oriented Directed Acyclic Graph
EARO:	Extended Address Registration Option
EDAC:	Extended Duplicate Address Confirmation
EDAR:	Extended Duplicate Address Request
ESS:	Extended Service Set
LLA:	Link-Layer Address
LLN:	Low-Power and Lossy Network
MLSN:	Multi-Link Subnet
MTU:	Maximum Transmission Unit
NA:	Neighbor Advertisement

NCE: Neighbor Cache Entry
ND: Neighbor Discovery
NS: Neighbor Solicitation
NUD: Neighbor Unreachability Detection
ODAD: Optimistic DAD
RA: Router Advertisement
ROVR: Registration Ownership Verifier
RPL: Routing Protocol for LLNs
RS: Router Solicitation
SLLAO: Source Link-Layer Address Option
SNMA: Solicited-Node Multicast Address
STA: Station
TID: Transaction ID
TLLAO: Target Link-Layer Address Option

2.4. Background

In this document, readers will encounter terms and concepts that are discussed in the following documents:

Classical IPv6 ND: "Neighbor Discovery for IP version 6 (IPv6)" [RFC4861], "IPv6 Stateless Address Autoconfiguration" [RFC4862], and "Optimistic Duplicate Address Detection (DAD) for IPv6" [RFC4429];

IPv6 ND over multiple links: "Neighbor Discovery Proxies (ND Proxy)" [RFC4389] and "Multi-Link Subnet Issues" [RFC4903];

6LoWPAN: "Problem Statement and Requirements for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing" [RFC6606]; and

6LoWPAN ND: Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs) [RFC6775], "Registration Extensions for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Neighbor Discovery" [RFC8505], and "Address-Protected Neighbor Discovery for Low-Power and Lossy Networks" [RFC8928].

3. Overview

This section and its subsections present a non-normative high-level view of the operation of the 6BBR. The following sections cover the normative part.

Figure 1 illustrates a Backbone Link that federates a collection of LLNs as a single IPv6 subnet, with a number of 6BBRs providing ND proxy services to their attached LLNs.

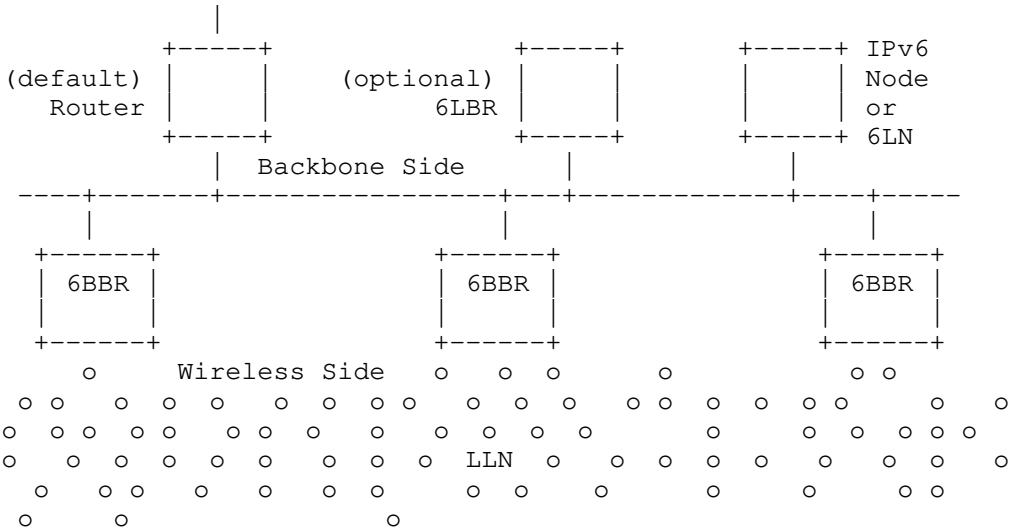

Figure 1: Backbone Link and Backbone Routers

The LLN may be a hub-and-spoke access link such as (Low-Power) IEEE Std 802.11 (Wi-Fi) [IEEEstd80211] and IEEE Std 802.15.1 (Bluetooth)

[IEEEstd802151] or a mesh-under or a route-over network [RFC8505]. The proxy state can be distributed across multiple 6BBRs attached to the same backbone.

The main features of a 6BBR are as follows:

* MLSN functions (provided by the 6BBR on the backbone) performed on behalf of Registered Nodes

* Routing Registrar services that reduce multicast within the LLN:

  - Binding Table management
  - failover, e.g., due to mobility

Each Backbone Router (6BBR) maintains a data structure for its Registered Addresses called a Binding Table. The abstract data that is stored in the Binding Table includes the Registered Address; anchor information on the Registering Node such as the connecting interface, link-local address, and link-layer address (LLA) of the Registering Node on that interface; the EARO including ROVR and TID; a state that can be either Reachable, Tentative, or Stale; and other information such as a trust level that may be configured, e.g., to protect a server. The combined Binding Tables of all the 6BBRs on a backbone form a distributed database of Registered Nodes that reside in the LLNs or on the IPv6 Backbone.

Unless otherwise configured, a 6BBR does the following:

* Creates a new entry in a Binding Table for a newly Registered Address and ensures that the address is not duplicated over the backbone.

* Advertises a Registered Address over the backbone using an NA message as either unsolicited or a response to an NS message. This includes joining the multicast group associated to the SNMA derived from the Registered Address, as specified in Section 7.2.1 of [RFC4861], over the backbone.

* The 6BBR MAY respond immediately as a proxy in lieu of the Registering Node, e.g., if the Registering Node has a sleep cycle that the 6BBR does not want to interrupt or if the 6BBR has a recent state that is deemed fresh enough to permit the proxied response. It is preferred, though, that the 6BBR checks whether the Registering Node is still responsive on the Registered Address. To that effect:

  - as a Bridging Proxy:
    the 6BBR forwards the multicast DAD and address lookup messages as a unicast link-layer frame to the link-layer address of the Registering Node that matches the target in the ND message; the Neighbor Unreachability Detection (NUD) message is unicast and is forwarded as is. In all cases, the goal is to let the Registering Node answer with the ND Message and options that it sees fit.
  - as a Routing Proxy:
    the 6BBR checks the liveliness of the Registering Node, e.g., using a NUD verification, before answering on its behalf.

* Delivers packets arriving from the LLN, using Neighbor Solicitation messages to look up the destination over the backbone.

* Forwards or bridges packets between the LLN and the backbone.

* Verifies liveness for a registration, when needed.

The first of these functions enables the 6BBR to fulfill its role as a Routing Registrar for each of its attached LLNs. The remaining functions fulfill the role of the 6BBRs as the border routers that federate the Multi-Link IPv6 Subnet.

The operation of IPv6 ND and ND proxy are not mutually exclusive on the backbone, meaning that nodes attached to the backbone and using IPv6 ND can transparently interact with 6LNs that rely on a 6BBR to ND proxy for them, whether the 6LNs are reachable over an LLN or directly attached to the backbone.

The registration mechanism [RFC8505] used to learn addresses to be proxied may coexist in a 6BBR with a proprietary snooping or the traditional bridging functionality of an AP, in order to support legacy LLN nodes that do not support this specification.

The registration to a proxy service uses an NS/NA exchange with EARO. The 6BBR operation resembles that of a Mobile IPv6 (MIPv6) [RFC6275] Home Agent (HA). The combination of a 6BBR and a MIPv6 HA enables full mobility support for 6LNs, inside and outside the links that form the subnet.

6BBRs perform IPv6 ND functions over the backbone as follows:

* The EARO [RFC8505] is used in IPv6 ND exchanges over the backbone between the 6BBRs to help distinguish duplication from movement. Extended Duplicate Address Messages (EDAR and EDAC) may also be used to communicate with a 6LBR, if one is present. Address duplication is detected using the ROVR field. Conflicting registrations to different 6BBRs for the same Registered Address are resolved using the TID field, which forms an order of registrations.

* The LLA that the 6BBR advertises for the Registered Address on behalf of the Registered Node over the backbone can belong to the Registering Node; in that case, the 6BBR (acting as a Bridging Proxy (see Section 8)) bridges the unicast packets. Alternatively, the LLA can be that of the 6BBR on the backbone interface, in which case, the 6BBR (acting as a Routing Proxy (see Section 7)) receives the unicast packets at Layer 3 and routes over.

3.1. Updating RFCs 6775 and 8505

This specification adds the EARO as a possible option in RS, NS(DAD), and NA messages over the backbone. This document specifies the use of those ND messages by 6BBRs over the backbone, at a high level in Section 6 and in more detail in Section 9.

> Note: [RFC8505] requires that the registration NS(EARO) contain a Source Link-Layer Address Option (SLLAO). [RFC4862] requires that the NS(DAD) be sent from the unspecified address for which there cannot be an SLLAO. Consequently, an NS(DAD) cannot be confused with a registration.

This specification allows the deployment of a 6LBR on the backbone where EDAR and EDAC messages coexist with classical ND. It also adds the capability to insert IPv6 ND options in the EDAR and EDAC messages. A 6BBR acting as a 6LR for the Registered Address can insert an SLLAO in the EDAR to the 6LBR in order to avoid causing a multicast NS(lookup) back. This enables the 6LBR to store the link-layer address associated with the Registered Address on a link and to serve as a mapping server as described in [UNICAST-LOOKUP].

This specification allows an address to be registered to more than one 6BBR. Consequently, a 6LBR that is deployed on the backbone MUST be capable of maintaining state for each of the 6BBRs that have registered with the same TID and same ROVR.

3.2. Access Link

The simplest MLSN topology from the Layer 3 perspective occurs when the wireless network appears as a single-hop hub-and-spoke network as shown in Figure 2. The Layer 2 operation may effectively be hub-and-spoke (e.g., Wi-Fi) or mesh-under, with a Layer 2 protocol handling the complex topology.

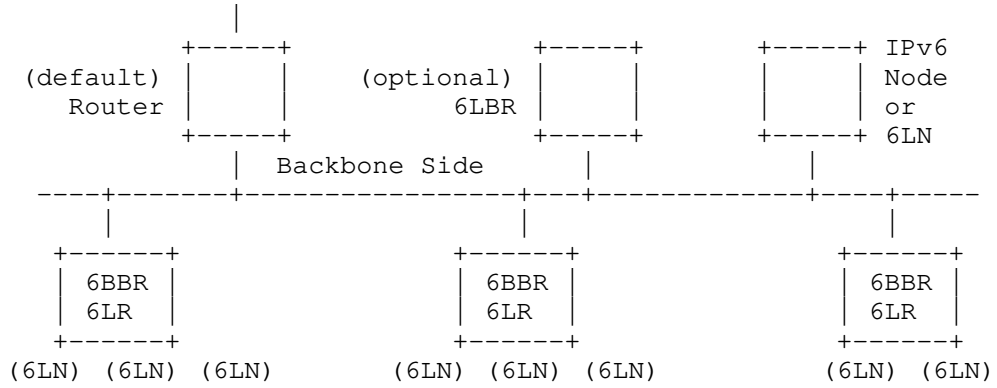

Figure 2: Access Link Use Case

Figure 3 illustrates a flow where 6LN forms an IPv6 address and registers it to a 6BBR acting as a 6LR [RFC8505]. The 6BBR applies Optimistic Duplicate Address Detection (ODAD) (see Section 3.6) to the Registered Address to enable connectivity while the message flow is still in progress.

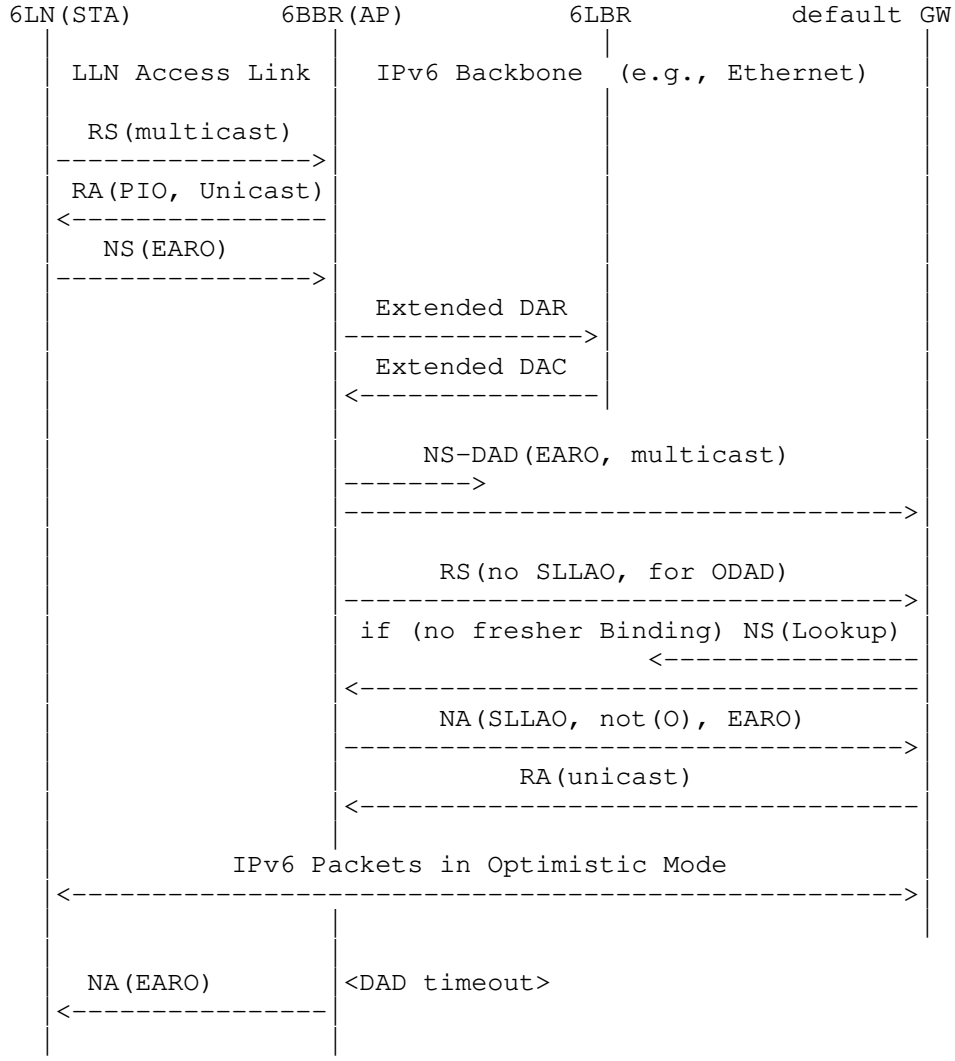

Figure 3: Initial Registration Flow to a 6BBR Acting as a Routing Proxy

In this example, a 6LBR is deployed on the Backbone Link to serve the whole subnet, and EDAR/EDAC messages are used in combination with DAD to enable coexistence with IPv6 ND over the backbone.

The RS sent initially by the 6LN (e.g., a Wi-Fi STA) is transmitted as a multicast, but since it is intercepted by the 6BBR, it is never effectively broadcast. The multiple arrows associated to the ND messages on the backbone denote a real Layer 2 broadcast.

3.3. Route-Over Mesh

A more complex MLSN topology occurs when the wireless network appears as a Layer 3 mesh network as shown in Figure 4. A so-called route-over routing protocol exposes routes between 6LRs towards both 6LRs and 6LNs, and a 6LBR acts as the Root of the Layer 3 mesh network and proxy-registers the LLN addresses to the 6BBR.

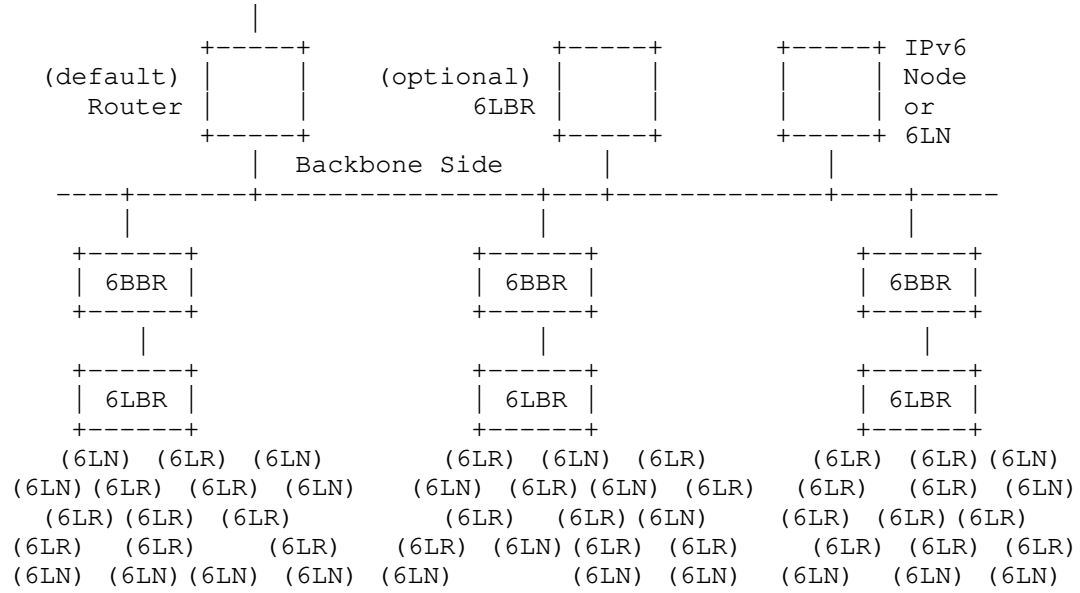

Figure 4: Route-Over Mesh Use Case

Figure 5 illustrates IPv6 signaling that enables a 6LN (the Registered Node) to form a Global or a Unique Local Address and register it to the 6LBR that serves its LLN using [RFC8505] and a neighboring 6LR as relay. The 6LBR (the Registering Node) then proxies the registration [RFC8505] to the 6BBR to obtain ND proxy services from the 6BBR.

The RS sent initially by the 6LN is transmitted as a multicast and contained within 1-hop broadcast range where hopefully a 6LR is found. The 6LR is expected to be already connected to the LLN and capable of reaching the 6LBR, which is possibly multiple hops away, using unicast messages.

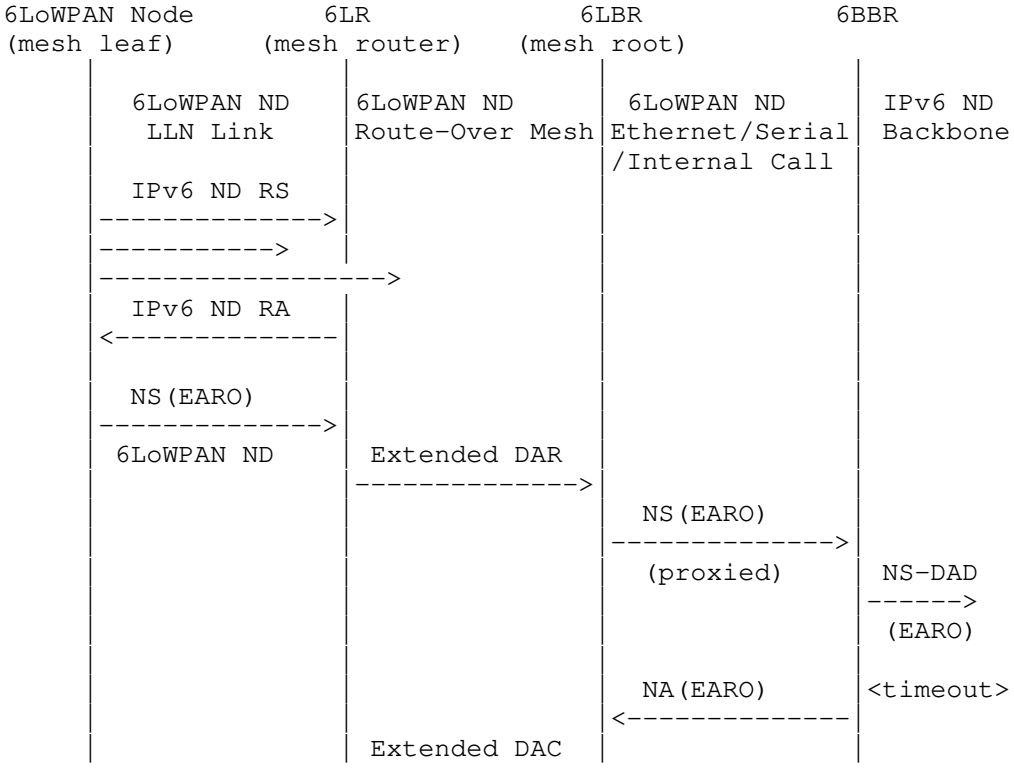

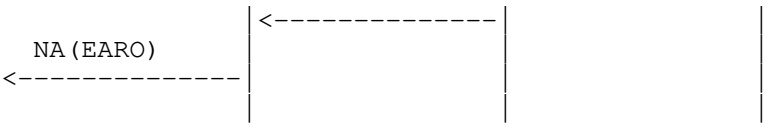

Figure 5: Initial Registration Flow over Route-Over Mesh

As a non-normative example of a route-over mesh, the IPv6 over the TSCH mode of IEEE 802.15.4e (6TiSCH) architecture [6TiSCH] suggests using the RPL [RFC6550] and collocating the RPL root with a 6LBR that serves the LLN. The 6LBR is also either collocated with or directly connected to the 6BBR over an IPv6 link.

3.4. The Binding Table

Addresses in an LLN that are reachable from the backbone by way of the 6BBR function must be registered to that 6BBR, using an NS(EARO) with the R flag set [RFC8505]. The 6BBR answers with an NA(EARO) and maintains a state for the registration in an abstract Binding Table.

An entry in the Binding Table is called a "Binding". A Binding may be in Tentative, Reachable, or Stale state.

The 6BBR uses a combination of [RFC8505] and IPv6 ND over the backbone to advertise the registration and avoid a duplication. Conflicting registrations are solved by the 6BBRs transparently to the Registering Nodes.

Only one 6LN may register a given address, but the address may be registered to multiple 6BBRs for higher availability.

Over the LLN, Binding Table management is as follows:

* De-registrations (newer TID, same ROVR, null Lifetime) are accepted with a status code of 4 ("Removed"); the entry is deleted.

* Newer registrations (newer TID, same ROVR, non-null Lifetime) are accepted with a status code of 0 ("Success"); the Binding is updated with the new TID, the Registration Lifetime, and the Registering Node. In Tentative state, the EDAC response is held and may be overwritten; in other states, the Registration Lifetime timer is restarted, and the entry is placed in Reachable state.

* Identical registrations (same TID, same ROVR) from the same Registering Node are accepted with a status code of 0 ("Success"). In Tentative state, the response is held and may be overwritten, but the response is eventually produced, carrying the result of the DAD process.

* Older registrations (older TID, same ROVR) from the same Registering Node are discarded.

* Identical and older registrations (not-newer TID, same ROVR) from a different Registering Node are rejected with a status code of 3 ("Moved"); this may be rate-limited to avoid undue interference.

* Any registration for the same address but with a different ROVR is rejected with a status code of 1 ("Duplicate Address").

The operation of the Binding Table is specified in detail in Section 9.

3.5. Primary and Secondary 6BBRs

A Registering Node MAY register the same address to more than one 6BBR, in which case, the Registering Node uses the same EARO in all the parallel registrations. On the other hand, there is no provision in 6LoWPAN ND for a 6LN (acting as Registered Node) to select its 6LBR (acting as Registering Node), so it cannot select more than one either. To allow for this, NS(DAD) and NA messages with an EARO

received over the backbone that indicate an identical Binding in another 6BBR (same Registered Address, same TID, same ROVR) are silently ignored except for the purpose of selecting the primary 6BBR for that registration.

A 6BBR may be either primary or secondary. The primary is the 6BBR that has the highest 64-bit Extended Unique Identifier (EUI-64) address of all the 6BBRs that share a registration for the same Registered Address, with the same ROVR and same Transaction ID, and the EUI-64 address is considered an unsigned 64-bit integer. A given 6BBR can be primary for a given address and secondary for another address, regardless of whether or not the addresses belong to the same 6LN.

In the following sections, it is expected that an NA will be sent over the backbone only if the node is primary or does not support the concept of primary. More than one 6BBR claiming or defending an address generates unwanted traffic, but there is no reachability issue since all 6BBRs provide reachability from the backbone to the 6LN.

If a Registering Node loses connectivity to its 6BBR or one of the 6BBRs to which it registered an address, it retries the registration to the (one or more) available 6BBR(s). When doing that, the Registering Node MUST increment the TID in order to force the migration of the state to the new 6BBR and the reselection of the primary 6BBR if it is the node that was lost.

3.6. Using Optimistic DAD

ODAD [RFC4429] specifies how an IPv6 address can be used before completion of DAD. ODAD guarantees that this behavior will not cause harm if the new address is a duplicate.

Support for ODAD avoids delays in installing the Neighbor Cache Entry (NCE) in the 6BBRs and the default router, enabling immediate connectivity to the Registered Node. As shown in Figure 3, if the 6BBR is aware of the LLA of a router, then the 6BBR sends a Router Solicitation (RS), using the Registered Address as the IP Source Address, to the known router(s). The RS is sent without an SLLAO, to avoid invalidating a preexisting NCE in the router.

Following ODAD, the router may then send a unicast RA to the Registered Address, and it may resolve that address using an NS(Lookup) message. In response, the 6BBR sends an NA with an EARO and the Override flag [RFC4861] that is not set. The router can then determine the freshest EARO in case of conflicting NA(EARO) messages, using the method described in Section 5.2.1 of [RFC8505]. If the NA(EARO) is the freshest answer, the default router creates a Binding with the SLLAO of the 6BBR (in Routing Proxy mode) or that of the Registering Node (in Bridging Proxy mode), so traffic from/to the Registered Address can flow immediately.

4. Multi-Link Subnet Considerations

The backbone and the federated LLN links are considered to be different links in the MLSN, even if multiple LLNs are attached to the same 6BBR. ND messages are link-scoped and are not forwarded by the 6BBR between the backbone and the LLNs, though some packets may be reinjected in Bridging Proxy mode (see Section 8).

Legacy nodes located on the backbone expect that the subnet is deployed within a single link and that there is a common Maximum Transmission Unit (MTU) for intra-subnet communication: the Link MTU. They will not perform the IPv6 Path MTU Discovery [RFC8201] for a destination within the subnet. For that reason, the MTU MUST have the same value on the backbone and on all federated LLNs in the MLSN. As a consequence, the 6BBR MUST use the same MTU value in RAs over the backbone and in the RAs that it transmits toward the LLN links.

5. Optional 6LBR Serving the Multi-Link Subnet

A 6LBR can be deployed to serve the whole MLSN as shown in Figure 4. It may be attached to the backbone, in which case it can be discovered by its capability advertisement (see Section 4.3 of [RFC8505]) in RA messages.

When a 6LBR is present, the 6BBR uses an EDAR/EDAC message exchange with the 6LBR to check if the new registration corresponds to a duplication or a movement. This is done prior to the NS(DAD) process, which may be avoided if the 6LBR already maintains a conflicting state for the Registered Address.

If this registration is a duplicate or not the freshest, then the 6LBR replies with an EDAC message with a status code of 1 ("Duplicate Address") or 3 ("Moved"), respectively. If this registration is the freshest, then the 6LBR replies with a status code of 0 ("Success"). In that case, if this registration is fresher than an existing registration for another 6BBR, then the 6LBR also sends an asynchronous EDAC with a status code of 4 ("Removed") to the older 6BBR.

The EDAR message SHOULD carry the SLLAO used in NS messages by the 6BBR for that Binding, and the EDAC message SHOULD carry the Target Link-Layer Address Option (TLLAO) associated with the currently accepted registration. This enables a 6BBR to locate the new position of a mobile 6LN in the case of a Routing Proxy operation and opens the capability for the 6LBR to serve as a mapping server in the future.

Note that if link-local addresses are registered, then the scope of uniqueness on which the address duplication is checked is the total collection of links that the 6LBR serves, as opposed to the sole link on which the link-local address is assigned.

6. Using IPv6 ND over the Backbone Link

On the backbone side, the 6BBR MUST join the SNMA group corresponding to a Registered Address as soon as it creates a Binding for that address and maintain that SNMA membership as long as it maintains the registration. The 6BBR uses either the SNMA or plain unicast to defend the Registered Addresses in its Binding Table over the backbone (as specified in [RFC4862]). The 6BBR advertises and defends the Registered Addresses over the Backbone Link using RS, NS(DAD), and NA messages with the Registered Address as the Source or Target Address.

The 6BBR MUST place an EARO in the IPv6 ND messages that it generates on behalf of the Registered Node. Note that an NS(DAD) does not contain an SLLAO and cannot be confused with a proxy registration such as performed by a 6LBR.

IPv6 ND operates as follows on the backbone:

* Section 7.2.8 of [RFC4861] specifies that an NA message generated as a proxy does not have the Override flag set in order to ensure that if the real owner is present on the link, its own NA will take precedence, and this NA does not update the NCE for the real owner if one exists.

* A node that receives multiple NA messages updates an existing NCE only if the Override flag is set; otherwise, the node will probe the cached address.

* When an NS(DAD) is received for a tentative address, which means that two nodes form the same address at nearly the same time, the node that first claimed the address cannot be detected per Section 5.4.3 of [RFC4862], and the address is abandoned.

* In any case, [RFC4862] indicates that a node never responds to a Neighbor Solicitation for a tentative address.

This specification adds information about proxied addresses that helps to sort out a duplication (different ROVR) from a movement (same ROVR, different TID); in the latter case, the older registration is sorted out from the fresher one (by comparing TIDs).

When a Registering Node moves from one 6BBR to the next, the 6BBRs send NA messages over the backbone to update existing NCEs. A node that receives multiple NA messages with an EARO option and the same ROVR MUST favor the NA with the freshest EARO over the others.

The new 6BBR MAY set the Override flag in the NA messages if it does not compete with the Registering Node for the NCE in backbone nodes. This is assured if the Registering Node is attached via an interface that cannot be bridged onto the backbone, making it impossible for the Registering Node to defend its own addresses there. This may also be signaled by the Registering Node through a protocol extension that is not in scope for this specification.

When the Binding is in Tentative state, the 6BBR acts as follows:

* an NS(DAD) that indicates a duplication can still not be asserted for first come, but the situation can be avoided using a 6LBR on the backbone that will serialize the order of appearance of the address and ensure first-come, first-served.

* an NS or an NA that denotes an older registration for the same Registered Node is not interpreted as a duplication as specified in Sections 5.4.3 and 5.4.4 of [RFC4862], respectively.

When the Binding is no longer in Tentative state, the 6BBR acts as follows:

* an NS or an NA with an EARO that denotes a duplicate registration (different ROVR) is answered with an NA message that carries an EARO with a status code of 1 ("Duplicate Address"), unless the received message is an NA that carries an EARO with a status code of 1 ("Duplicate Address").

In any state, the 6BBR acts as follows:

* an NS or an NA with an EARO that denotes an older registration (same ROVR) is answered with an NA message that carries an EARO with a status code of 3 ("Moved") to ensure that the Stale state is removed rapidly.

This behavior is specified in more detail in Section 9.

This specification enables proxy operation for the IPv6 ND resolution of LLN devices, and a prefix that is used across an MLSN MAY be advertised as on-link over the backbone. This is done for backward compatibility with existing IPv6 hosts by setting the L flag in the Prefix Information Option (PIO) of RA messages [RFC4861].

For movement involving a slow reattachment, the NUD procedure defined in [RFC4861] may timeout too quickly. Nodes on the backbone SHOULD support [RFC7048] whenever possible.

7. Routing Proxy Operations

A Routing Proxy provides IPv6 ND proxy functions for Global and Unique Local Addresses between the LLN and the backbone, but not for link-local addresses. It operates as an IPv6 border router and provides a full link-layer isolation.

In this mode, it is not required that the link-layer addresses of the 6LNs be visible at Layer 2 over the backbone. Thus, it is useful when the messaging over the backbone that is associated with wireless mobility becomes expensive, e.g., when the Layer 2 topology is virtualized over a wide area IP underlay.

This mode is definitely required when the LLN uses a link-layer

address format that is different from that on the backbone (e.g., EUI-64 versus EUI-48). Since a 6LN may not be able to resolve an arbitrary destination in the MLSN directly, a prefix that is used across a MLSN MUST NOT be advertised as on-link in RA messages sent towards the LLN.

In order to maintain IP connectivity, the 6BBR installs a connected host route to the Registered Address on the LLN interface, via the Registering Node as identified by the source address and the SLLAO in the NS(EARO) messages.

When operating as a Routing Proxy, the 6BBR MUST use its Layer 2 address on its backbone interface in the SLLAO of the RS messages and the TLLAO of the NA messages that it generates to advertise the Registered Addresses.

For each Registered Address, multiple peers on the backbone may have resolved the address with the 6BBR link-layer address, maintaining that mapping in their Neighbor Cache. The 6BBR SHOULD maintain a list of the peers on the backbone that have associated its link-layer address with the Registered Address. If that Registered Address moves to another 6BBR, the previous 6BBR SHOULD unicast a gratuitous NA to each such peer, to supply the LLA of the new 6BBR in the TLLAO for the address. A 6BBR that does not maintain this list MAY multicast a gratuitous NA message; this NA will possibly hit all the nodes on the backbone, whether or not they maintain an NCE for the Registered Address. In either case, the 6BBR MAY set the Override flag if it is known that the Registered Node cannot attach to the backbone; this will avoid interruptions and save probing flows in the future.

If a correspondent fails to receive the gratuitous NA, it will keep sending traffic to a 6BBR to which the node was previously registered. Since the previous 6BBR removed its host route to the Registered Address, it will look up the address over the backbone, resolve the address with the LLA of the new 6BBR, and forward the packet to the correct 6BBR. The previous 6BBR SHOULD also issue a redirect message [RFC4861] to update the cache of the correspondent.

8. Bridging Proxy Operations

A Bridging Proxy provides IPv6 ND proxy functions between the LLN and the backbone while preserving the forwarding continuity at the link layer. It acts as a Layer 2 bridge for all types of unicast packets including link-scoped, and it appears as an IPv6 Host on the backbone.

The Bridging Proxy registers any Binding, including a link-local address to the 6LBR (if present), and defends it over the backbone in IPv6 ND procedures.

To achieve this, the Bridging Proxy intercepts the IPv6 ND messages and may reinject them on the other side, respond directly, or drop them. For instance, an NS(Lookup) from the backbone that matches a Binding can be responded to directly or turned into a unicast on the LLN side to let the 6LN respond.

As a Bridging Proxy, the 6BBR MUST use the Registering Node's Layer 2 address in the SLLAO of the NS/RS messages and the TLLAO of the NA messages that it generates to advertise the Registered Addresses. The Registering Node's Layer 2 address is found in the SLLAO of the registration NS(EARO) and maintained in the Binding Table.

The MLSN prefix SHOULD NOT be advertised as on-link in RA messages sent towards the LLN. If a destination address is seen as on-link, then a 6LN may use NS(Lookup) messages to resolve that address. In that case, the 6BBR MUST either answer the NS(Lookup) message directly or reinject the message on the backbone, as either a Layer 2 unicast or a multicast.

If the Registering Node owns the Registered Address, meaning that the

Registering Node is the Registered Node, then its mobility does not impact existing NCEs over the backbone. In a network where proxy registrations are used, meaning that the Registering Node acts on behalf of the Registered Node, if the Registered Node selects a new Registering Node, then the existing NCEs across the backbone pointing at the old Registering Node must be updated. In that case, the 6BBR SHOULD attempt to fix the existing NCEs across the backbone pointing at other 6BBRs using NA messages as described in Section 7.

This method can fail if the multicast message is not received; one or more correspondent nodes on the backbone might maintain a stale NCE, and packets to the Registered Address may be lost. When this condition happens, it is eventually discovered and resolved using NUD as defined in [RFC4861].

9. Creating and Maintaining a Binding

Upon receiving a registration for a new address (i.e., an NS(EARO) with the R flag set), the 6BBR creates a Binding and operates as a 6LR according to [RFC8505], interacting with the 6LBR if one is present.

An implementation of a Routing Proxy that creates a Binding MUST also create an associated host route pointing to the Registering Node in the LLN interface from which the registration was received.

Acting as a 6BBR, the 6LR operation is modified as follows:

* Acting as a Bridging Proxy, the 6LR MUST ND proxy over the backbone for registered link-local addresses.

* EDAR and EDAC messages SHOULD carry an SLLAO and a TLLAO, respectively.

* An EDAC message with a status code of 9 ("6LBR Registry Saturated") is assimilated as a status code of 0 ("Success") if a following DAD process protects the address against duplication.

This specification enables nodes on a Backbone Link to coexist along with nodes implementing IPv6 ND [RFC4861] as well as other non-normative specifications such as [SAVI-WLAN]. It is possible that not all IPv6 addresses on the backbone are registered and known to the 6LBR, and an EDAR/EDAC exchange with the 6LBR might succeed even for a duplicate address. Consequently, the 6BBR still needs to perform IPv6 ND DAD over the backbone after an EDAC with a status code of 0 ("Success") or 9 ("6LBR Registry Saturated").

For the DAD operation, the Binding is placed in Tentative state for a duration of TENTATIVE_DURATION (Section 12), and an NS(DAD) message is sent as a multicast message over the backbone to the SNMA associated with the Registered Address [RFC4862]. The EARO from the registration MUST be placed unchanged in the NS(DAD) message.

If a registration is received for an existing Binding with a non-null Registration Lifetime and the registration is fresher (same ROVR, fresher TID), then the Binding is updated with the new Registration Lifetime, TID, and possibly Registering Node. In Tentative state (see Section 9.1), the current DAD operation continues unaltered. In other states (see Sections 9.2 and 9.3), the Binding is placed in Reachable state for the Registration Lifetime, and the 6BBR returns an NA(EARO) to the Registering Node with a status code of 0 ("Success").

Upon a registration that is identical (same ROVR, TID, and Registering Node), the 6BBR does not alter its current state. In Reachable state, it returns an NA(EARO) back to the Registering Node with a status code of 0 ("Success"). A registration that is not as fresh (same ROVR, older TID) is ignored.

If a registration is received for an existing Binding and a Registration Lifetime of 0, then the Binding is removed, and the 6BBR

returns an NA(EARO) back to the Registering Node with a status code of 0 ("Success"). An implementation of a Routing Proxy that removes a Binding MUST remove the associated host route pointing on the Registering Node.

The old 6BBR removes its Binding Table entry and notifies the Registering Node with a status code of 3 ("Moved") if a new 6BBR claims a fresher registration (same ROVR, fresher TID) for the same address. The old 6BBR MAY preserve a temporary state in order to forward packets in flight. The state may be, for instance, an NCE that was formed when an NA message was received. It may also be a Binding Table entry in Stale state, pointing at the new 6BBR on the backbone or any other abstract cache entry that can be used to resolve the link-layer address of the new 6BBR. The old 6BBR SHOULD also use REDIRECT messages pointing at the new 6BBR to update the correspondents of the Registered Address, as specified in [RFC4861].

9.1. Operations on a Binding in Tentative State

The Tentative state covers a DAD period over the backbone during which an address being registered is checked for duplication using the procedures defined in [RFC4862].

For a Binding in Tentative state:

* The Binding MUST be removed if an NA message is received over the backbone for the Registered Address with no EARO or with an EARO that indicates an existing registration owned by a different Registering Node (different ROVR). In that case, an NA is sent back to the Registering Node with a status code of 1 ("Duplicate Address") to indicate that the Binding has been rejected. This behavior might be overridden by policy, in particular if the registration is trusted, e.g., based on the validation of the ROVR field (see [RFC8928]).

* The Binding MUST be removed if an NS(DAD) message is received over the backbone for the Registered Address with no EARO or with an EARO that has a different ROVR that indicates a tentative registration by a different Registering Node. In that case, an NA is sent back to the Registering Node with a status code of 1 ("Duplicate Address"). This behavior might be overridden by policy, in particular if the registration is trusted, e.g., based on the validation of the ROVR field (see [RFC8928]).

* The Binding MUST be removed if an NA or an NS(DAD) message is received over the backbone for the Registered Address and contains an EARO that indicates a fresher registration [RFC8505] for the same Registering Node (same ROVR). In that case, an NA MUST be sent back to the Registering Node with a status code of 3 ("Moved").

* The Binding MUST be kept unchanged if an NA or an NS(DAD) message is received over the backbone for the Registered Address and contains an EARO that indicates an older registration [RFC8505] for the same Registering Node (same ROVR). The message is answered with an NA that carries an EARO with a status code of 3 ("Moved") and the Override flag not set. This behavior might be overridden by policy, in particular if the registration is not trusted.

* Other NS(DAD) and NA messages from the backbone are ignored.

* NS(Lookup) and NS(NUD) messages SHOULD be optimistically answered with an NA message containing an EARO with a status code of 0 ("Success") and the Override flag not set (see Section 3.6). If optimistic DAD is disabled, then they SHOULD be queued to be answered when the Binding goes to Reachable state.

When the TENTATIVE_DURATION (Section 12) timer elapses, the Binding is placed in Reachable state for the Registration Lifetime, and the 6BBR returns an NA(EARO) to the Registering Node with a status code

of 0 ("Success").

The 6BBR also attempts to take over any existing Binding from other 6BBRs and to update existing NCEs in backbone nodes. This is done by sending an NA message with an EARO and the Override flag not set over the backbone (see Sections 7 and 8).

9.2. Operations on a Binding in Reachable State

The Reachable state covers an active registration after a successful DAD process.

If the Registration Lifetime is of a long duration, an implementation might be configured to reassess the availability of the Registering Node at a lower period, using a NUD procedure as specified in [RFC7048]. If the NUD procedure fails, the Binding SHOULD be placed in Stale state immediately.

For a Binding in Reachable state:

* The Binding MUST be removed if an NA or an NS(DAD) message is received over the backbone for the Registered Address and contains an EARO that indicates a fresher registration [RFC8505] for the same Registered Node (i.e., same ROVR but fresher TID). A status code of 4 ("Removed") is returned in an asynchronous NA(EARO) to the Registering Node. Based on configuration, an implementation may delay this operation by a timer with a short setting, e.g., a few seconds to a minute, in order to allow for a parallel registration to reach this node, in which case the NA might be ignored.

* NS(DAD) and NA messages containing an EARO that indicates a registration for the same Registered Node that is not as fresh as this Binding MUST be answered with an NA message containing an EARO with a status code of 3 ("Moved").

* An NS(DAD) with no EARO or with an EARO that indicates a duplicate registration (i.e., different ROVR) MUST be answered with an NA message containing an EARO with a status code of 1 ("Duplicate Address") and the Override flag not set, unless the received message is an NA that carries an EARO with a status code of 1 ("Duplicate Address"), in which case the node refrains from answering.

* Other NS(DAD) and NA messages from the backbone are ignored.

* NS(Lookup) and NS(NUD) messages SHOULD be answered with an NA message containing an EARO with a status code of 0 ("Success") and the Override flag not set. The 6BBR MAY check whether the Registering Node is still available using a NUD procedure over the LLN prior to answering; this behavior depends on the use case and is subject to configuration.

When the Registration Lifetime timer elapses, the Binding is placed in Stale state for a duration of STALE_DURATION (Section 12).

9.3. Operations on a Binding in Stale State

The Stale state enables tracking of the backbone peers that have a NCE pointing to this 6BBR in case the Registered Address shows up later.

If the Registered Address is claimed by another 6LN on the backbone, with an NS(DAD) or an NA, the 6BBR does not defend the address.

For a Binding in Stale state:

* The Binding MUST be removed if an NA or an NS(DAD) message is received over the backbone for the Registered Address with no EARO or with an EARO that indicates either a fresher registration for the same Registered Node or a duplicate registration. A status

code of 4 ("Removed") MAY be returned in an asynchronous NA(EARO) to the Registering Node.

* NS(DAD) and NA messages containing an EARO that indicates a registration for the same Registered Node that is not as fresh as this MUST be answered with an NA message containing an EARO with a status code of 3 ("Moved").

* If the 6BBR receives an NS(Lookup) or an NS(NUD) message for the Registered Address, the 6BBR MUST attempt a NUD procedure as specified in [RFC7048] to the Registering Node, targeting the Registered Address, prior to answering. If the NUD procedure succeeds, the operation in Reachable state applies. If the NUD fails, the 6BBR refrains from answering.

* Other NS(DAD) and NA messages from the backbone are ignored.

When the STALE_DURATION (Section 12) timer elapses, the Binding MUST be removed.

10. Registering Node Considerations

A Registering Node MUST implement [RFC8505] in order to interact with a 6BBR (which acts as a Routing Registrar). Following [RFC8505], the Registering Node signals that it requires IPv6 ND proxy services from a 6BBR by registering the corresponding IPv6 address using an NS(EARO) message with the R flag set.

The Registering Node may be the 6LN owning the IPv6 address or a 6LBR that performs the registration on its behalf in a route-over mesh.

A 6LN MUST register all of its IPv6 addresses to its 6LR, which is the 6BBR when they are connected at Layer 2. Failure to register an address may result in the address being unreachable by other parties. This would happen, for instance, if the 6BBR propagates the NS(Lookup) from the backbone only to the LLN nodes that do not register their addresses.

The Registering Node MUST refrain from using multicast NS(Lookup) when the destination is not known as on-link, e.g., if the prefix is advertised in a PIO with the L flag not set. In that case, the Registering Node sends its packets directly to its 6LR.

The Registering Node SHOULD also follow BCP 202 [RFC7772] in order to limit the use of multicast RAs. It SHOULD also implement "Simple Procedures for Detecting Network Attachment in IPv6" [RFC6059] (DNA procedures) to detect movements and support "Packet-Loss Resiliency for Router Solicitations" [RFC7559] in order to improve reliability for the unicast RS messages.

11. Security Considerations

The procedures in this document modify the mechanisms used for IPv6 ND and DAD and should not affect other aspects of IPv6 or higher-level-protocol operation. As such, the main classes of attacks that are in play are those that work to block Neighbor Discovery or to forcibly claim an address that another node is attempting to use. In the absence of cryptographic protection at higher layers, the latter class of attacks can have significant consequences, with the attacker being able to read all the "stolen" traffic that was directed to the target of the attack.

This specification applies to LLNs and a backbone in which the individual links are protected against rogue access on the LLN by authenticating a node that attaches to the network and encrypting the transmissions at the link layer and on the backbone side, using the physical security and access control measures that are typically applied there; thus, packets may neither be forged nor overheard.

In particular, the LLN link layer is required to provide secure unicast to/from the Backbone Router and secure broadcast from the

routers in a way that prevents tampering with or replaying the ND messages.

For the IPv6 ND operation over the backbone, and unless the classical ND is disabled (e.g., by configuration), the classical ND messages are interpreted as emitted by the address owner and have precedence over the 6BBR that is only a proxy.

As a result, the security threats that are detailed in Section 11.1 of [RFC4861] fully apply to this specification as well. In short:

* Any node that can send a packet on the backbone can take over any address, including addresses of LLN nodes, by claiming it with an NA message and the Override bit set. This means that the real owner will stop receiving its packets.

* Any node that can send a packet on the backbone can forge traffic and pretend it is issued from an address that it does not own, even if it did not claim the address using ND.

* Any node that can send a packet on the backbone can present itself as a preferred router to intercept all traffic outgoing on the subnet. It may even expose a prefix on the subnet as "not-on-link" and intercept all the traffic within the subnet.

* If the rogue can receive a packet from the backbone, it can also snoop all the intercepted traffic, by stealing an address or the role of a router.

This means that any rogue access to the backbone must be prevented at all times, and nodes that are attached to the backbone must be fully trusted / never compromised.

Using address registration as the sole ND mechanism on a link and coupling it with [RFC8928] guarantees the ownership of a Registered Address within that link.

* The protection is based on a proof of ownership encoded in the ROVR field, and it protects against address theft and impersonation by a 6LN, because the 6LR can challenge the Registered Node for a proof of ownership.

* The protection extends to the full LLN in the case of an LLN link, but it does not extend over the backbone since the 6BBR cannot provide the proof of ownership when it defends the address.

A possible attack over the backbone can be done by sending an NS with an EARO and expecting the NA(EARO) back to contain the TID and ROVR fields of the existing state. With that information, the attacker can easily increase the TID and take over the Binding.

If the classical ND is disabled on the backbone and the use of [RFC8928] and a 6LBR are mandated, the network will benefit from the following new advantages:

Zero-trust security for ND flows within the whole subnet: the increased security that [RFC8928] provides on the LLN will also apply to the backbone; it becomes impossible for an attached node to claim an address that belongs to another node using ND, and the network can filter packets that are not originated by the owner of the source address (Source Address Validation Improvement (SAVI)), as long as the routers are known and trusted.

Remote ND DoS attack avoidance: the complete list of addresses in the network will be known to the 6LBR and available to the default router; with that information, the router does not need to send a multicast NS(Lookup) in case of a Neighbor Cache miss for an incoming packet, which is a source of remote DoS attack against the network.

Less IPv6 ND-related multicast on the backbone: DAD and NS(Lookup)

become unicast queries to the 6LBR.

Better DAD operation on wireless: DAD has been found to fail to detect duplications on large Wi-Fi infrastructures due to the unreliable broadcast operation on wireless; using a 6LBR enables a unicast lookup.

Less Layer 2 churn on the backbone: Using the Routing Proxy approach, the link-layer address of the LLN devices and their mobility are not visible in the backbone; only the link-Layer addresses of the 6BBR and backbone nodes are visible at Layer 2 on the backbone. This is mandatory for LLNs that cannot be bridged on the backbone and useful in any case to scale down, stabilize the forwarding tables at Layer 2, and avoid the gratuitous frames that are typically broadcasted to fix the transparent bridging tables when a wireless node roams from an AP to the next.

This specification introduces a 6BBR that is a router on the path of the LLN traffic and a 6LBR that is used for the lookup. They could be interesting targets for an attacker. A compromised 6BBR can accept a registration but block the traffic or refrain from proxying. A compromised 6LBR may unduly accept the transfer of ownership of an address or block a newcomer by faking that its address is a duplicate. But those attacks are possible in a classical network from a compromised default router and a DHCP server, respectively, and can be prevented using the same methods.

A possible attack over the LLN can still be done by compromising a 6LR. A compromised 6LR may modify the ROVR of EDAR messages in flight and transfer the ownership of the Registered Address to itself or a tier. It may also claim that a ROVR was validated when it really wasn't and reattribute an address to itself or to an attached 6LN. This means that 6LRs, as well as 6LBRs and 6BBRS, must still be fully trusted / never compromised.

This specification mandates checking on the 6LBR on the backbone before doing the classical DAD, in case the address already exists. This may delay the DAD operation and should be protected by a short timer, in the order of 100 ms or less, which will only represent a small extra delay versus the 1 s wait of the DAD operation.

12. Protocol Constants

This specification uses the following constants:

TENTATIVE_DURATION: 800 milliseconds

In LLNs with long-lived addresses such as Low-Power WAN (LPWANs), STALE_DURATION SHOULD be configured with a relatively long value to cover an interval when the address may be reused and before it is safe to expect that the address was definitively released. A good default value is 24 hours. In LLNs where addresses are renewed rapidly, e.g., for privacy reasons, STALE_DURATION SHOULD be configured with a relatively shorter value -- 5 minutes by default.

13. IANA Considerations

This document has no IANA actions.

14. 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>.

[RFC6059] Krishnan, S. and G. Daley, "Simple Procedures for Detecting Network Attachment in IPv6", RFC 6059, DOI 10.17487/RFC6059, November 2010, <https://www.rfc-editor.org/info/rfc6059>.

[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>.

[RFC7048] Nordmark, E. and I. Gashinsky, "Neighbor Unreachability Detection Is Too Impatient", RFC 7048, DOI 10.17487/RFC7048, January 2014, <https://www.rfc-editor.org/info/rfc7048>.

[RFC7559] Krishnan, S., Anipko, D., and D. Thaler, "Packet-Loss Resiliency for Router Solicitations", RFC 7559, DOI 10.17487/RFC7559, May 2015, <https://www.rfc-editor.org/info/rfc7559>.

[RFC7772] Yourtchenko, A. and L. Colitti, "Reducing Energy Consumption of Router Advertisements", BCP 202, RFC 7772, DOI 10.17487/RFC7772, February 2016, <https://www.rfc-editor.org/info/rfc7772>.

[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>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

[RFC8201] McCann, J., Deering, S., Mogul, J., and R. Hinden, Ed., "Path MTU Discovery for IP version 6", STD 87, RFC 8201, DOI 10.17487/RFC8201, July 2017, <https://www.rfc-editor.org/info/rfc8201>.

[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>.

15. Informative References

[6TiSCH] Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", Work in Progress, Internet-Draft, draft-ietf-6tisch-architecture-29, 27 August 2020, <https://tools.ietf.org/html/draft-ietf-6tisch-architecture-29>.

[DAD-APPROACHES] Nordmark, E., "Possible approaches to make DAD more robust and/or efficient", Work in Progress, Internet-Draft, draft-nordmark-6man-dad-approaches-02, 19 October 2015,

<https://tools.ietf.org/html/draft-nordmark-6man-dad-approaches-02>.

[DAD-ISSUES]
Yourtchenko, A. and E. Nordmark, "A survey of issues related to IPv6 Duplicate Address Detection", Work in Progress, Internet-Draft, draft-yourtchenko-6man-dad-issues-01, 3 March 2015, <https://tools.ietf.org/html/draft-yourtchenko-6man-dad-issues-01>.

[IEEEstd80211]
IEEE, "IEEE Standard for Information technology--Telecommunications and information exchange between systems Local and metropolitan area networks--Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", IEEE 802.11-2012, DOI 10.1109/ieeestd.2016.7786995, December 2016, <https://ieeexplore.ieee.org/document/7786995>.

[IEEEstd802151]
IEEE, "IEEE Standard for Information technology--Local and metropolitan area networks--Specific requirements--Part 15.1a: Wireless Medium Access Control (MAC) and Physical Layer (PHY) specifications for Wireless Personal Area Networks (WPAN)", IEEE 802.15.1-2005, DOI 10.1109/ieeestd.2005.96290, June 2005, <https://ieeexplore.ieee.org/document/1490827>.

[IEEEstd802154]
IEEE, "IEEE Standard for Local and metropolitan area networks--Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs)", IEEE 802.15.4-2011, DOI 10.1109/ieeestd.2011.6012487, September 2011, <https://ieeexplore.ieee.org/document/6012487>.

[IEEEstd8021Q]
IEEE, "IEEE Standard for Local and Metropolitan Area Networks--Bridges and Bridged Networks", IEEE 802.1Q-2018, DOI 10.1109/IEEESTD.2018.8403927, July 2018, <https://ieeexplore.ieee.org/document/8403927>.

[MCAST-PROBLEMS]
Perkins, C. E., McBride, M., Stanley, D., Kumari, W., and J. C. Zuniga, "Multicast Considerations over IEEE 802 Wireless Media", Work in Progress, Internet-Draft, draft-ietf-mboned-ieee802-mcast-problems-12, 26 October 2020, <https://tools.ietf.org/html/draft-ietf-mboned-ieee802-mcast-problems-12>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <https://www.rfc-editor.org/info/rfc4271>.

[RFC4389] Thaler, D., Talwar, M., and C. Patel, "Neighbor Discovery Proxies (ND Proxy)", RFC 4389, DOI 10.17487/RFC4389, April 2006, <https://www.rfc-editor.org/info/rfc4389>.

[RFC4903] Thaler, D., "Multi-Link Subnet Issues", RFC 4903, DOI 10.17487/RFC4903, June 2007, <https://www.rfc-editor.org/info/rfc4903>.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008, <https://www.rfc-editor.org/info/rfc5340>.

[RFC5415] Calhoun, P., Ed., Montemurro, M., Ed., and D. Stanley, Ed., "Control And Provisioning of Wireless Access Points (CAPWAP) Protocol Specification", RFC 5415, DOI 10.17487/RFC5415, March 2009,

<https://www.rfc-editor.org/info/rfc5415>.

[RFC6275] Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, DOI 10.17487/RFC6275, July 2011, <https://www.rfc-editor.org/info/rfc6275>.

[RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <https://www.rfc-editor.org/info/rfc6550>.

[RFC6606] Kim, E., Kaspar, D., Gomez, C., and C. Bormann, "Problem Statement and Requirements for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing", RFC 6606, DOI 10.17487/RFC6606, May 2012, <https://www.rfc-editor.org/info/rfc6606>.

[RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", RFC 6830, DOI 10.17487/RFC6830, January 2013, <https://www.rfc-editor.org/info/rfc6830>.

[RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", RFC 7432, DOI 10.17487/RFC7432, February 2015, <https://www.rfc-editor.org/info/rfc7432>.

[RFC8273] Brzozowski, J. and G. Van de Velde, "Unique IPv6 Prefix per Host", RFC 8273, DOI 10.17487/RFC8273, December 2017, <https://www.rfc-editor.org/info/rfc8273>.

[RFC8928] Thubert, P., Ed., Sarikaya, B., Sethi, M., and R. Struik, "Address-Protected Neighbor Discovery for Low-Power and Lossy Networks", RFC 8928, DOI 10.17487/RFC8928, November 2020, <https://www.rfc-editor.org/info/rfc8928>.

[RIFT] Przygienda, T., Sharma, A., Thubert, P., Rijsman, B., and D. Afanasiev, "RIFT: Routing in Fat Trees", Work in Progress, Internet-Draft, draft-ietf-rift-rift-12, 26 May 2020, <https://tools.ietf.org/html/draft-ietf-rift-rift-12>.

[RPL-LEAVES] Thubert, P. and M. C. Richardson, "Routing for RPL Leaves", Work in Progress, Internet-Draft, draft-ietf-roll-unaware-leaves-23, 10 November 2020, <https://tools.ietf.org/html/draft-ietf-roll-unaware-leaves-23>.

[RS-REFRESH] Nordmark, E., Yourtchenko, A., and S. Krishnan, "IPv6 Neighbor Discovery Optional RS/RA Refresh", Work in Progress, Internet-Draft, draft-ietf-6man-rs-refresh-02, 31 October 2016, <https://tools.ietf.org/html/draft-ietf-6man-rs-refresh-02>.

[SAVI-WLAN] Bi, J., Wu, J., Wang, Y., and T. Lin, "A SAVI Solution for WLAN", Work in Progress, Internet-Draft, draft-bi-savi-wlan-20, 14 November 2020, <https://tools.ietf.org/html/draft-bi-savi-wlan-20>.

[UNICAST-LOOKUP] Thubert, P. and E. Levy-Abegnoli, "IPv6 Neighbor Discovery Unicast Lookup", Work in Progress, Internet-Draft, draft-thubert-6lo-unicast-lookup-00, 25 January 2019, <https://tools.ietf.org/html/draft-thubert-6lo-unicast-lookup-00>.

Appendix A. Possible Future Extensions

With the current specification, the 6LBR is not leveraged to avoid multicast NS(Lookup) on the backbone. This could be done by adding a lookup procedure in the EDAR/EDAC exchange.

By default, the specification does not have a fine-grained trust model: all nodes that can authenticate to the LLN link layer or attach to the backbone are equally trusted. It would be desirable to provide a stronger authorization model, e.g., whereby nodes that associate their address with a proof of ownership [RFC8928] should be trusted more than nodes that do not. Such a trust model and related signaling could be added in the future to override the default operation and favor trusted nodes.

As an alternate to the ND Proxy operation, the registration may be redistributed as a host route in a routing protocol that would operate over the backbone; this is already happening in IoT networks [RPL-LEAVES] and Data Center Routing [RIFT] and could be extended to other protocols, e.g., BGP [RFC4271] and OSPFv3 [RFC5340]. The registration may also be advertised in an overlay protocol such as Mobile IPv6 (MIPv6) [RFC6275], the Locator/ID Separation Protocol (LISP) [RFC6830], or Ethernet VPN (EVPN) [RFC7432].

Appendix B. Applicability and Requirements Served

This document specifies ND proxy functions that can be used to federate an IPv6 Backbone Link and multiple IPv6 LLNs into a single MLSN. The ND proxy functions enable IPv6 ND services for DAD and address lookup that do not require broadcasts over the LLNs.

The term LLN is used to cover multiple types of WLANs and WPANs, including (Low-Power) Wi-Fi, BLUETOOTH(R) Low Energy, IEEE Std 802.11ah and IEEE Std 802.15.4 wireless meshes, and the types of networks listed in "Requirements Related to Various Low-Power Link Types" (see Appendix B.3 of [RFC8505]).

Each LLN in the subnet is attached to a 6BBR. The Backbone Routers interconnect the LLNs and advertise the addresses of the 6LNs over the Backbone Link using ND proxy operations.

This specification updates IPv6 ND over the backbone to distinguish address movement from duplication and eliminate Stale state in the backbone routers and backbone nodes once a 6LN has roamed. This way, mobile nodes may roam rapidly from one 6BBR to the next, and requirements are met per "Requirements Related to Mobility" (see Appendix B.1 of [RFC8505]).

A 6LN can register its IPv6 addresses and thereby obtain ND proxy services over the backbone, meeting the requirements expressed in "Requirements Related to Proxy Operations" (see Appendix B.4 of [RFC8505].

The negative impact of the IPv6 ND-related broadcasts can be limited to one of the federated links, enabling the number of 6LNs to grow. The Routing Proxy operation avoids the need to expose the link-layer addresses of the 6LNs onto the backbone, keeping the Layer 2 topology simple and stable. This meets the requirements in "Requirements Related to Scalability" (see Appendix B.6 of [RFC8505]), as long as the 6BBRs are dimensioned for the number of registrations that each needs to support.

In the case of a Wi-Fi access link, a 6BBR may be collocated with the AP, a Fabric Edge (FE), or a Control and Provisioning of Wireless Access Points (CAPWAP) [RFC5415] Wireless LAN Controller (WLC). In those cases, the wireless client (STA) is the 6LN that makes use of [RFC8505] to register its IPv6 address(es) to the 6BBR acting as the Routing Registrar. The 6LBR can be centralized and either connected to the Backbone Link or reachable over IP. The 6BBR ND proxy operations eliminate the need for wireless nodes to respond synchronously when a lookup is performed for their IPv6 addresses.

This provides the function of a Sleep Proxy for ND [DAD-APPROACHES].

For the Time-Slotted Channel Hopping (TSCH) mode of [IEEEstd802154], the 6TiSCH architecture [6TiSCH] describes how a 6LoWPAN ND host could connect to the Internet via a RPL mesh network, but doing so requires extensions to the 6LOWPAN ND protocol to support mobility and reachability in a secure and manageable environment. The extensions detailed in this document also work for the 6TiSCH architecture, serving the requirements listed in "Requirements Related to Routing Protocols" (see Appendix B.2 of [RFC8505]).

The registration mechanism may be seen as a more reliable alternate to snooping [SAVI-WLAN]. Note that registration and snooping are not mutually exclusive. Snooping may be used in conjunction with the registration for nodes that do not register their IPv6 addresses. The 6BBR assumes that if a node registers at least one IPv6 address to it, then the node registers all of its addresses to the 6BBR. With this assumption, the 6BBR can possibly cancel all undesirable multicast NS messages that would otherwise have been delivered to that node.

Scalability of the MLSN [RFC4903] requires avoidance of multicast/ broadcast operations as much as possible even on the backbone [MCAST-PROBLEMS]. Although hosts can connect to the backbone using IPv6 ND operations, multicast RAs can be saved by using [RS-REFRESH], which also requires the support of [RFC7559].

Acknowledgments

Many thanks to Dorothy Stanley, Thomas Watteyne, and Jerome Henry for their various contributions. Also, many thanks to Timothy Winters and Erik Nordmark for their help, review, and support in preparation for the IESG cycle and to Kyle Rose, Elwyn Davies, Barry Leiba, Mirja KÃ¼hlewind, Alvaro Retana, Roman Danyliw, and especially Dominique Barthel and Benjamin Kaduk for their useful contributions through the IETF Last Call and IESG process.

Authors' Addresses

Pascal Thubert (editor)
Cisco Systems, Inc.
Building D
45 Allee des Ormes - BP1200
06254 MOUGINS - Sophia Antipolis
France

Phone: +33 497 23 26 34
Email: pthubert@cisco.com

Charles E. Perkins
Blue Meadow Networking
Saratoga, CA 95070
United States of America

Email: charliep@computer.org

Eric Levy-Abegnoli
Cisco Systems, Inc.
Building D
45 Allee des Ormes - BP1200
06254 MOUGINS - Sophia Antipolis
France

Phone: +33 497 23 26 20
Email: elevyabe@cisco.com