Workgroup: 6lo
Updates: [6775](), [8505]() (if approved)
Published: 3 March 2020
Intended Status: Standards Track
Expires: 4 September 2020

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

IPv6 Backbone Router

Abstract

This document updates RFC 6775 and RFC 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.

Status of This Memo

This Internet-Draft is submitted in full conformance with the
provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering
Task Force (IETF). Note that other groups may also distribute
working documents as Internet-Drafts. The list of current Internet-
Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six
months and may be updated, replaced, or obsoleted by other documents
at any time. It is inappropriate to use Internet-Drafts as reference
material or to cite them other than as "work in progress."

This Internet-Draft will expire on 4 September 2020.

Copyright Notice

Section 4.e of the Trust Legal Provisions and are provided without
warranty as described in the Simplified BSD License.

Table of Contents

1. **Introduction**

[IEEE STD. 802.1](#) [[IEEEstd8021](#)] Ethernet Bridging 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
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 behaviours when
employed unmodified on a local wireless medium (see [[I-D.ietf-mboned-ieee802-mcast-problems](#)]).

[Wi-Fi](#) [[IEEEstd80211](#)] Access Points (APs) deployed in an Extended
Service Set (ESS) act as Ethernet Bridges [[IEEEstd8021](#)], with the
property that the bridging state is established at the time of
association. This ensures connectivity to the end node (the Wi-Fi
STA) and protects the wireless medium against broadcast-intensive
Transparent Bridging reactive Lookups. In other words, the
association process is used to register the MAC 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, IPv6 [[RFC8200](#)] Neighbor
Discovery [[RFC4861](#)] [[RFC4862](#)] Protocol (IPv6 ND) 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 [I-D.yourtchenko-6man-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 succesful operation of the IPv6 ND duplicate address detection 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, and 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 and/or a LowPower Lossy Network (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 MAC 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 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 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, 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, 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 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 proxy-ND services. A 6BBR acting as a Bridging Proxy provides a proxy-ND function with Layer-2 continuity and can be collocated with a Wi-Fi Access Point (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 for can be obtained by snooping the IPV6 ND protocol (see [I-D.bi-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 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 Neighbor Discovery (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 to it as unicast the corresponding multicast packets.

2. Terminology

2.1. BCP 14

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 Multi-
   Link Subnet. The proxy-ND operation of 6BBRs over the Backbone
   extends IPv6 ND operation over the access links.

Sleeping Proxy: A 6BBR acts as a Sleeping Proxy if it answers IPv6
   ND Neighbor Solicitations over the Backbone on behalf of the
   Registering Node that is in a sleep state and 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 MAC
   Address as the Target Link Layer Address (TLLA) in the proxied
   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 MAC Layer. In that
   case, the MAC Address and the mobility of the Registering Node is
   visible across the bridged Backbone. The Bridging Proxy
   advertises the MAC Address of the Registering Node as the TLLA in
   the proxied NAs over the Backbone, and 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 NUD messages that target the Registered
   Address to the Registering Node as unicast frames and let it
   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 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

ARO: Address Registration Option

DAC: Duplicate Address Confirmation

DAD: Duplicate Address Detection

DAR: Duplicate Address Request

EARO: Extended Address Registration Option

EDAC: Extended Duplicate Address Confirmation

EDAR: Extended Duplicate Address Request

DODAG: Destination-Oriented Directed Acyclic Graph

ID: Identifier

LLN: Low-Power and Lossy Network

NA: Neighbor Advertisement

MAC: Medium Access Control

NCE: Neighbor Cache Entry

ND: Neighbor Discovery

NDP: Neighbor Discovery Protocol

NS: Neighbor Solicitation

NS(DAD): NDP NS message used for the purpose of duplication
   avoidance (multicast)

NS(Lookup): NDP NS message used for the purpose of address
   resolution (multicast)

NS(NUD): NDP NS message used for the purpose of unreachability
   detection (unicast)

NUD: Neighbor Unreachability Detection

ROVR: Registration Ownership Verifier

RPL: IPv6 Routing Protocol for LLNs

RA: Router Advertisement

RS: Router Solicitation

SNMA: Solicited-Node Multicast Address

LLA: Link Layer Address (aka MAC address)

SLLA: Source Link Layer Address

TLLA: Target Link Layer Address

TID: Transaction ID

2.4. References

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" [RFC4861],
"IPv6 Stateless Address Autoconfiguration" [RFC4862] and
"Optimistic Duplicate Address Detection" [RFC4429],

IPv6 ND over multiple links: "Neighbor Discovery Proxies (proxy-ND)" [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 Low-Power and Lossy Networks [RFC6775] and "Registration Extensions for 6LoWPAN Neighbor Discovery" [RFC8505].

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 proxy-ND services to their attached LLNs.

```
              |
          +-----+                +-----+        +-----+ IPv6
 (default) |     |   (Optional) |     |        |     | Node
   Router  |     |         6LBR |     |        |     | or
          +-----+                +-----+        +-----+ 6LN
           |  Backbone side       |              |
  ----+-------+----------------+---+------------+----+-----
      |                        |              |
  +------+                 +------+           +------+
  | 6BBR |                 | 6BBR |           | 6BBR |
  |      |                 |      |           |      |
  +------+                 +------+           +------+
      o    Wireless side  o  o o      o         o o
   o o   o o  o   o o  oo  o  o o   o o o  o  o o     o   o
  o  o o  o o   o o  o   o   o  o  o o        o     o  o  o o o
  o   o o o  o  o   o  o  o LLN o  o  o o  o   o   o  o  o  o
   o   o o  o   o   o o    o  o   o     o     o    o o
   o      o                 o
```
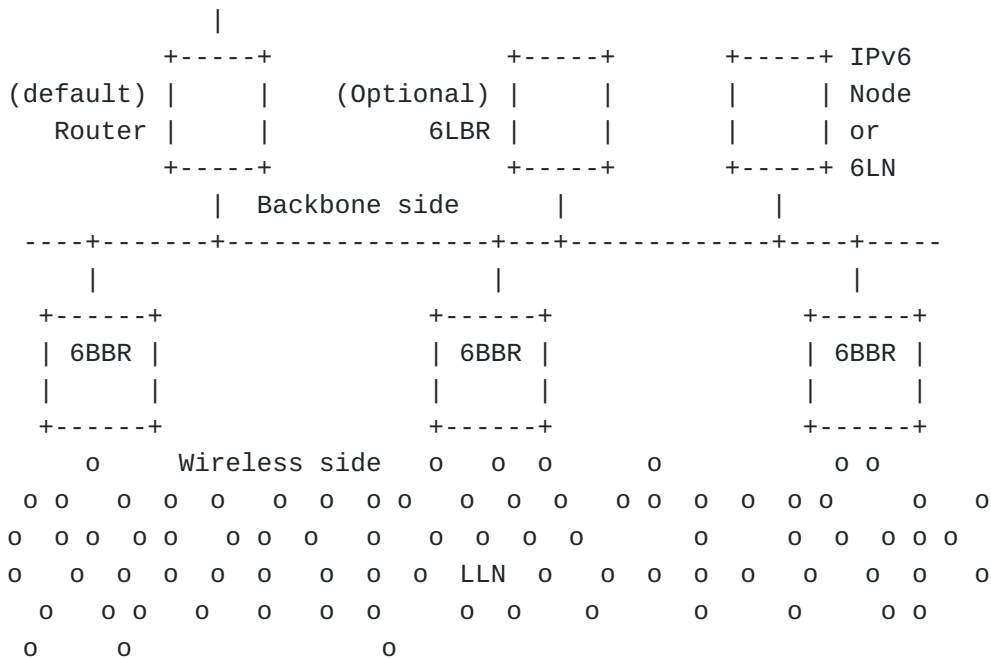
Figure 1: Backbone Link and Backbone Routers

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:

  *Multi-Link-subnet functions (provided by the 6BBR on the
   backbone) performed on behalf of Registered Nodes, and

  *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 connecting
interface, Link-Local Address and Link-Layer Address 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:

  *Create a new entry in a Binding Table for a new Registered
   Address and ensure that the Address is not duplicated over the
   Backbone.

  *Advertise a Registered Address over the Backbone using an NA
   message, either unsolicited or as a response to a 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 sleeping
   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 MAC-Layer frames to the MAC address of
     the Registering Node that matches the Target in the ND

message, and forwards as is the unicast Neighbor
Unreachability Detection (NUD) messages, so as 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.

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

  *Forward or bridge packets between the LLN and the Backbone.

  *Verify 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 of proxy-ND 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 proxy ND for them, whether the 6LNs are reachable over an
LLN or directly attached to the Backbone.

The [RFC8505] registration mechanism used to learn addresses to be
proxied may co-exist in a 6BBR with a proprietary snooping or the
traditional bridging functionality of an Access Point, 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.

The 6BBRs performs IPv6 ND functions over the backbone as follows:

  *The EARO [RFC8505] is used in the 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 Link Layer Address (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 RFC 6775 and RFC 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) contains an
Source Link Layer Address Option (SLLAO). [RFC4862] requires that
the NS(DAD) is sent from the unspecified address for which there
cannot be a SLLAO. Consequently, an NS(DAD) cannot be confused with
a registration.

This specification allows to deploy 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 a Lookup back. This
enables the 6LBR to store the MAC address associated to the
Registered Address on a Link and to serve as a mapping server as
described in [I-D.thubert-6lo-unicast-lookup].

This specification allows for 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 6BBR having
registered with the same TID and same ROVR.

3.2. Access Link

The simplest Multi-Link Subnet 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.

```
                   |
             +-----+              +-----+       +-----+ IPv6
   (default) |     |  (Optional)  |     |       |     | Node
      Router |     |       6LBR   |     |       |     | or
             +-----+              +-----+       +-----+ 6LN
             |    Backbone side        |             |
    ----+-------+-----------------+---+-------------+----+-----
        |                         |                 |
     +------+                  +------+           +------+
     | 6BBR |                  | 6BBR |           | 6BBR |
     | 6LR  |                  | 6LR  |           | 6LR  |
     +------+                  +------+           +------+
   (6LN) (6LN) (6LN)        (6LN) (6LN) (6LN)    (6LN) (6LN)
```
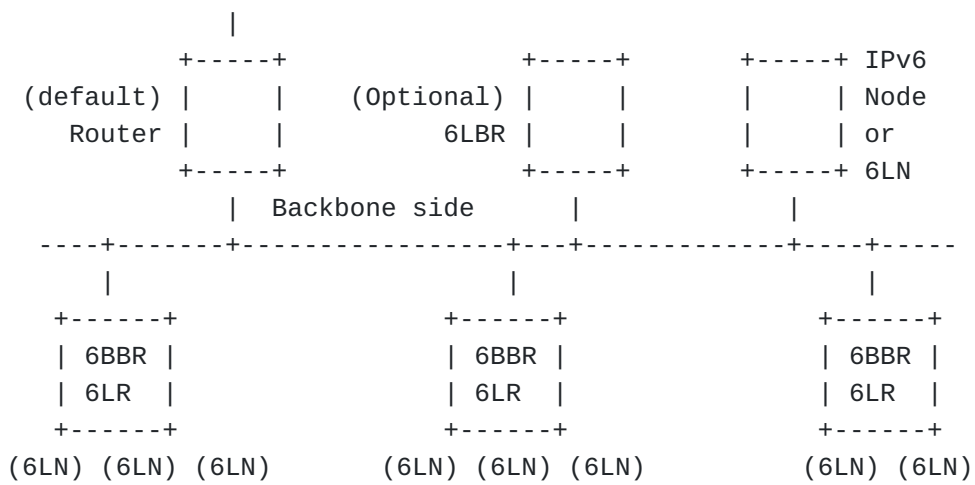
                  Figure 2: Access Link Use case

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
ODAD (see Section 3.6) to the registered address to enable
connectivity while the message flow is still in progress.

```
      6LN(STA)              6BBR(AP)              6LBR            default GW
         |                     |                     |                     |
         | LLN Access Link |  IPv6 Backbone  (e.g., Ethernet)   |
         |                     |                     |                     |
         |   RS(multicast)  |                     |                     |
         |---------------->|                     |                     |
         | RA(PIO, Unicast)|                     |                     |
         |<----------------|                     |                     |
         |    NS(EARO)      |                     |                     |
         |---------------->|                     |                     |
         |                     |  Extended DAR  |                     |
         |                     |--------------->|                     |
         |                     |  Extended DAC  |                     |
         |                     |<---------------|                     |
         |                     |                     |                     |
         |                     |    NS-DAD(EARO, multicast)      |
         |                     |-------->                         |
         |                     |---------------------------------->|
         |                     |                     |                     |
         |                     |      RS(no SLLAO, for ODAD)      |
         |                     |---------------------------------->|
         |                     | if (no fresher Binding) NS(Lookup) |
         |                     |                  <---------------|
         |                     |<----------------------------------|
         |                     |    NA(SLLAO, not(O), EARO)      |
         |                     |---------------------------------->|
         |                     |           RA(unicast)           |
         |                     |<----------------------------------|
         |                     |                     |                     |
         |         IPv6 Packets in optimistic mode        |
         |<------------------------------------------------------>|
         |                     |                     |                     |
         |                     |                     |                     |
         |   NA(EARO)      |<DAD timeout>     |                     |
         |<----------------|                     |                     |
         |                     |                     |                     |
```
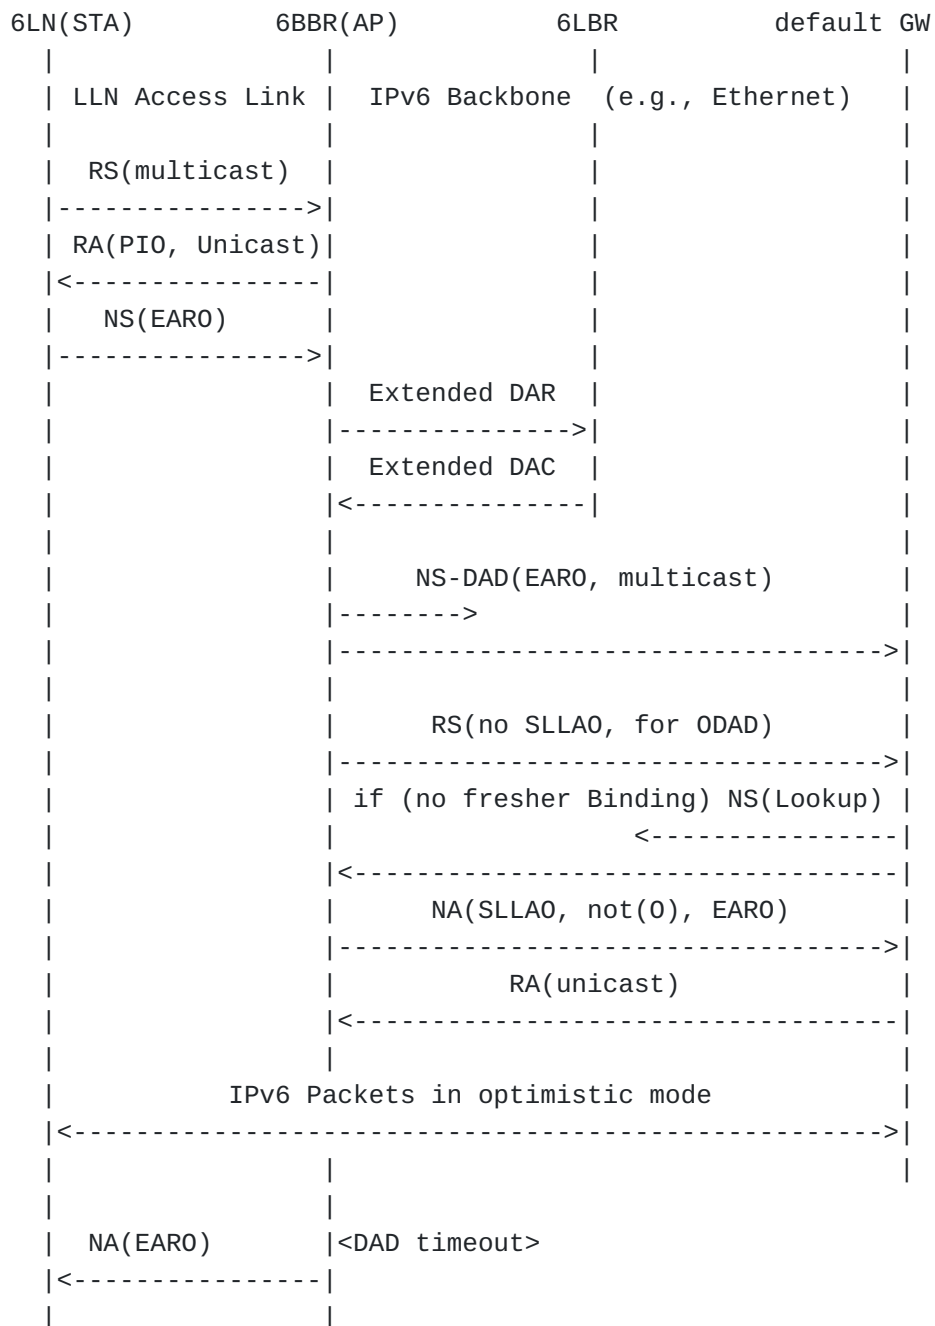
       Figure 3: Initial Registration Flow to a 6BBR acting as Routing Proxy

   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 Multi-Link Subnet 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 Root of the Layer 3
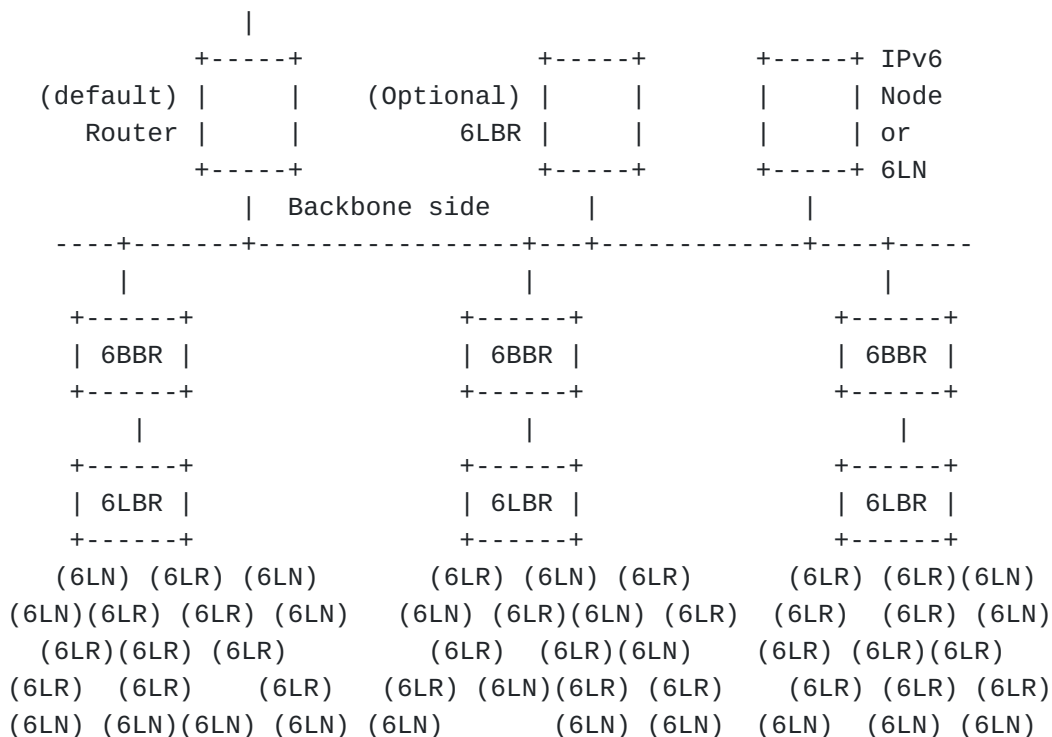Mesh network and proxy-registers the LLN addresses to the 6BBR.

```
                 |
          +-----+                 +-----+          +-----+ IPv6
  (default) |     |   (Optional) |     |          |     | Node
    Router |     |        6LBR  |     |          |     | or
          +-----+                 +-----+          +-----+ 6LN
          |  Backbone side       |                 |
  ----+-------+-----------------+---+-------------+----+-----
     |                     |                 |
   +------+                 +------+              +------+
   | 6BBR |                 | 6BBR |              | 6BBR |
   +------+                 +------+              +------+
      |                        |                    |
   +------+                 +------+              +------+
   | 6LBR |                 | 6LBR |              | 6LBR |
   +------+                 +------+              +------+
   (6LN) (6LR) (6LN)       (6LR) (6LN) (6LR)      (6LR) (6LR)(6LN)
 (6LN)(6LR) (6LR) (6LN)    (6LN) (6LR)(6LN) (6LR)  (6LR)  (6LR) (6LN)
  (6LR)(6LR) (6LR)         (6LR)  (6LR)(6LN)     (6LR) (6LR)(6LR)
 (6LR)  (6LR)   (6LR)     (6LR) (6LN)(6LR) (6LR)   (6LR) (6LR) (6LR)
 (6LN) (6LN)(6LN) (6LN) (6LN)        (6LN) (6LN)  (6LN)  (6LN) (6LN)
```

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] using a
neighboring 6LR as relay. The 6LBR (the Registering Node) then
proxies the [RFC8505] registration to the 6BBR to obtain proxy-ND
services from the 6BBR.

The RS sent initially by the 6LN is a 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
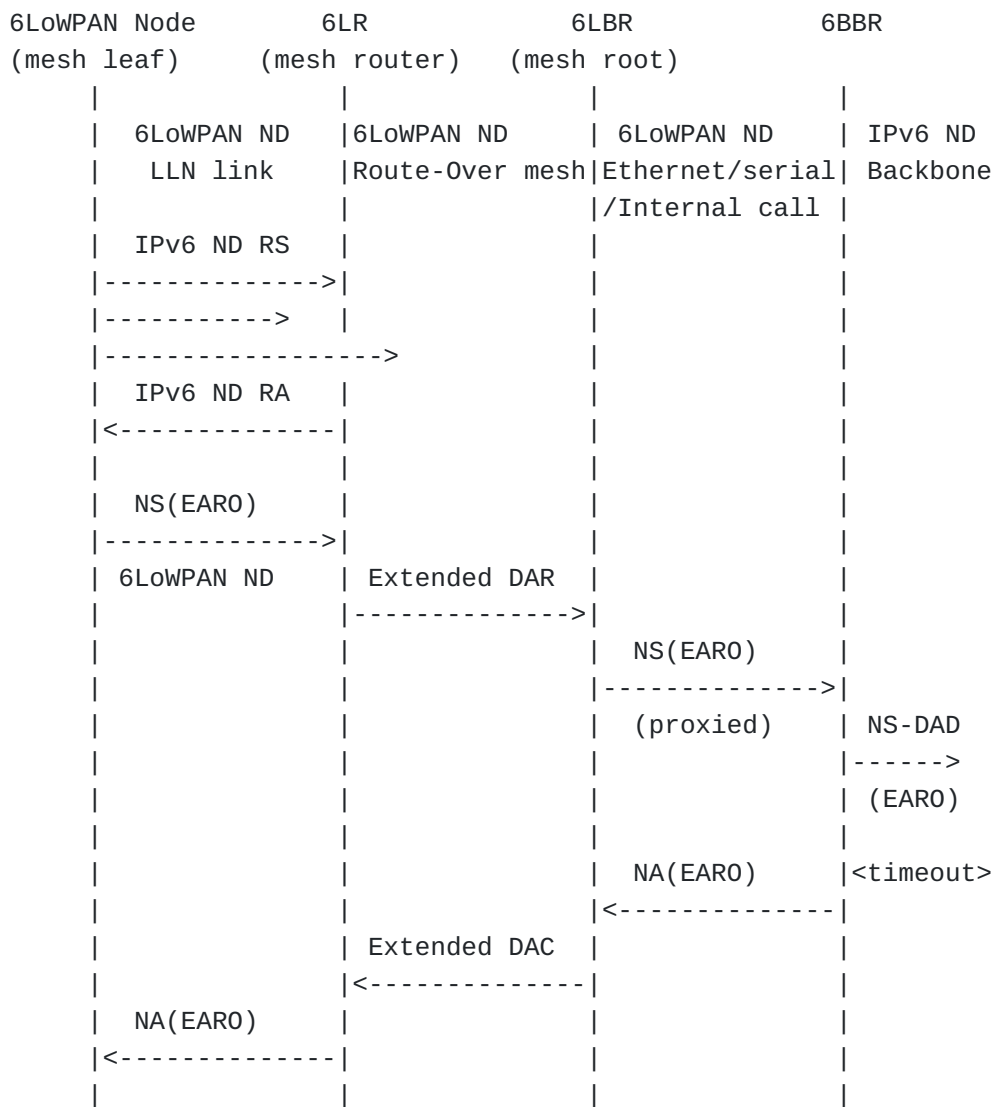capable to reach the 6LBR, possibly multiple hops away, using
unicast messages.

```
6LoWPAN Node         6LR                6LBR              6BBR
(mesh leaf)     (mesh router)   (mesh root)
     |               |               |                  |
     |  6LoWPAN ND   |6LoWPAN ND     | 6LoWPAN ND       | IPv6 ND
     |   LLN link    |Route-Over mesh|Ethernet/serial|  Backbone
     |               |               |/Internal call |  |
     |  IPv6 ND RS   |               |               |  |
     |-------------->|               |               |  |
     |---------->    |               |               |  |
     |------------------>            |               |  |
     |  IPv6 ND RA   |               |               |  |
     |<--------------|               |               |  |
     |               |               |               |  |
     |   NS(EARO)    |               |               |  |
     |-------------->|               |               |  |
     | 6LoWPAN ND    | Extended DAR  |               |  |
     |               |-------------->|               |  |
     |               |               |   NS(EARO)    |  |
     |               |               |-------------->|  |
     |               |               |  (proxied)    | NS-DAD
     |               |               |               |  |------>
     |               |               |               |  | (EARO)
     |               |               |               |  |
     |               |               |   NA(EARO)    |<timeout>
     |               |               |<--------------|  |
     |               | Extended DAC  |               |  |
     |               |<--------------|               |  |
     |   NA(EARO)    |               |               |  |
     |<--------------|               |               |  |
     |               |               |               |  |
```

                Figure 5: Initial Registration Flow over Route-Over Mesh

   As a non-normative example of a Route-Over Mesh, the 6TiSCH
   architecture [I-D.ietf-6tisch-architecture] suggests using the RPL
   [RFC6550] routing protocol 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 of 4 ("Removed"); the entry is deleted;

   *Newer registrations (newer TID, same ROVR, non-null Lifetime) are
    accepted with a status 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 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 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 of 1 (Duplicate).

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 but 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 EUI-64 Address of all the 6BBRs that share a
registration for the same Registered Address, with the same ROVR and
same Transaction ID, the EUI-64 Address being considered as an
unsigned 64bit 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 is 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 no reachability issue since
all 6BBRs provide reachability from the Backbone to the 6LN.

If a Registering Node loses connectivity to its 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

Optimistic Duplicate Address Detection [RFC4429] (ODAD) specifies
how an IPv6 Address can be used before completion of Duplicate
Address Detection (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 Link-Layer Address (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 a Source LLA Option (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 that traffic
from/to the Registered Address can flow immediately.

4. Multi-Link Subnet Considerations

The Backbone and the federated LLN Links are considered as different
links in the Multi-Link Subnet, 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 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 towards the LLN links.

5. Optional 6LBR serving the Multi-Link Subnet

A 6LBR can be deployed to serve the whole MLSN. 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 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. 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 of 4 ("Removed") to that other 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 that 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, section 5.4.3 of [RFC4862] cannot detect which node first claimed the address 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 sort out a duplication (different ROVR) from a movement (same ROVR, different TID), and in the latter case the older registration from the fresher one (by comparing TIDs).

When a Registering Node moves from one 6BBR to the next, the new 6BBR sends NA messages over the backbone to update existing NCEs. A node that supports this specification and 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 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-serve.

  *an NS or an NA that denotes an older registration for the same
   Registered Node is not interpreted as a duplication as specified
   in section 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 of 1 (Duplicate), unless the received message
   is an NA that carries an EARO with a status of 1.

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 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 a Multi-
Link Subnet 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 time out 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 MAC addresses of the 6LNs
are visible at Layer 2 over the Backbone. It is thus useful when the
messaging over the Backbone that is associated to 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 MAC address
format that is different from that on the Backbone (e.g., EUI-64 vs.
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 SLLA
option 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 MAC Address, maintaining that
mapping in their Neighbor Cache. The 6BBR SHOULD maintain a list of
the peers on the Backbone which have associated its MAC 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 TLLA option 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, so as
to 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
   MAC Layer. It acts as a Layer 2 Bridge for all types of unicast
   packets including link-scoped, and appears as an IPv6 Host on the
   Backbone.

   The Bridging Proxy registers any Binding including for 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 ND(Lookup) from the backbone that matches a
   Binding can be responded 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 SLLA of the
   registration NS(EARO), and maintained in the Binding Table.

   The Multi-Link Subnet 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, either as
   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 an 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 Bridging Proxy the 6LR MUST proxy ND over the backbone for registered Link-Local Addresses.

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

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

This specification enables nodes on a Backbone Link to co-exist along with nodes implementing IPv6 ND [RFC4861] as well as other non-normative specifications such as [I-D.bi-savi-wlan]. It is possible that not all IPv6 addresses on the Backbone are registered and known to the 6LBR, and an EDAR/EDAC echange 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 or 9.

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 Section 9.2 and Section 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 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 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 zero, then the Binding is removed, and the 6BBR returns an NA(EARO) back to the Registering Node with a status 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 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 for instance be a NCE formed based on a received NA message. It may also be a Binding Table entry in Stale state and 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 as specified in [RFC4861] to update the correspondents for the Registered Address, pointing to the new 6BBR.

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 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 containing 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 of 1 (Duplicate) 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 [I-D.ietf-6lo-ap-nd]).

  *The Binding MUST be removed if an NS(DAD) message is received over the Backbone for the Registered Address with no EARO, or containing an EARO with 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 of 1 (Duplicate). 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 [I-D.ietf-6lo-ap-nd]).

*The Binding MUST be removed if an NA or an NS(DAD) message is
   received over the Backbone for the Registered Address containing
   an EARO with a 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 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
   containing an EARO with a 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
   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 of 0 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 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 Section 7 and Section 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 containing
   an EARO that indicates a fresher registration ([RFC8505]) for the

same Registered Node (i.e., same ROVR but fresher TID). A status
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 a 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 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 of 1
 (Duplicate) and the Override flag not set, unless the received
 message is an NA that carries an EARO with a status of 1, 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 of 0 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 behaviour 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 containing
 no EARO or an EARO that indicates either a fresher registration
 for the same Registered Node or a duplicate registration. A
 status 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 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 proxy-
ND 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 that is 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

This specification applies to LLNs and a backbone in which the
individual links are protected against rogue access, by
authenticating a node that attaches to the network and encrypting at
the MAC layer the transmissions so packets may neither be overheard
or nor forged. In particular, the LLN MAC 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. It results that the
security threats that are detailed in section 11.1 of [RFC4861]
fully apply to this specification as well. In very 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 a 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
   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, be it 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 that 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 [I-D.ietf-6lo-ap-nd] 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 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 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 [I-D.ietf-6lo-ap-nd] 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 [I-D.ietf-6lo-ap-nd] 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 (SAVI), as long as that 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 NA(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 introduce 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, but not more than a default router and a DHCP server, respectively, which already exist in classical networks, and can be defended 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 self 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 to check 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 100ms or less, which will only represent a small extra delay versus the 1s wait of the DAD operation.

12. Protocol Constants

This Specification uses the following constants:

TENTATIVE_DURATION: 800 milliseconds

STALE_DURATION: see below

In LLNs with long-lived Addresses such as 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 can be 24 hours. In LLNs where addresses are renewed rapidly, e.g., for privacy reasons, STALE_DURATION SHOULD be configured with a relatively shorter value, by default 5 minutes.

13. IANA Considerations

This document has no request to IANA.

14. 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 to the IESG cycle, and to Kyle Rose, Elwyn Davies, Barry Leiba, Mirja Kühlewind, Alvaro Retana, Roman Danyliw and very especially Dominique Barthel and Benjamin Kaduk for their useful contributions through the IETF last call and IESG process.

15. 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>.

16. Informative References

[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>.

[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>.

[RFC5568] Koodli, R., Ed., "Mobile IPv6 Fast Handovers", RFC 5568,
           DOI 10.17487/RFC5568, July 2009, <https://www.rfc-editor.org/info/rfc5568>.

[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>.

[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>.

[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>.

[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>.

[I-D.yourtchenko-6man-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>.

[I-D.nordmark-6man-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>.

[I-D.ietf-6man-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>.

[I-D.ietf-6lo-ap-nd]
           Thubert, P., Sarikaya, B., Sethi, M., and R. Struik, "Address Protected Neighbor Discovery for Low-power and Lossy Networks", Work in Progress, Internet-Draft, draft-ietf-6lo-ap-nd-19, 6 February 2020, <https://tools.ietf.org/html/draft-ietf-6lo-ap-nd-19>.

[I-D.ietf-6tisch-architecture]
      Thubert, P., "An Architecture for IPv6 over the TSCH mode
      of IEEE 802.15.4", Work in Progress, Internet-Draft,
      draft-ietf-6tisch-architecture-28, 29 October 2019,
      <https://tools.ietf.org/html/draft-ietf-6tisch-
      architecture-28>.

[I-D.ietf-mboned-ieee802-mcast-problems]
      Perkins, C., McBride, M., Stanley, D., Kumari, W., and J.
      Zuniga, "Multicast Considerations over IEEE 802 Wireless
      Media", Work in Progress, Internet-Draft, draft-ietf-
      mboned-ieee802-mcast-problems-11, 11 December 2019,
      <https://tools.ietf.org/html/draft-ietf-mboned-ieee802-
      mcast-problems-11>.

[I-D.bi-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-18, 17 November 2019, <https://
      tools.ietf.org/html/draft-bi-savi-wlan-18>.

[I-D.thubert-6lo-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>.

[IEEEstd8021] IEEE standard for Information Technology, "IEEE
      Standard for Information technology -- Telecommunications
      and information exchange between systems Local and
      metropolitan area networks Part 1: Bridging and
      Architecture", .

[IEEEstd80211] IEEE standard for Information Technology, "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", .

[IEEEstd802151] IEEE standard for Information Technology, "IEEE
      Standard for Information Technology - Telecommunications
      and Information Exchange Between Systems - Local and
      Metropolitan Area Networks - Specific Requirements. -
      Part 15.1: Wireless Medium Access Control (MAC) and

Physical Layer (PHY) Specifications for Wireless Personal
Area Networks (WPANs)", .

[IEEEstd802154] IEEE standard for Information Technology, "IEEE
Standard for Local and metropolitan area networks -- Part
15.4: Low-Rate Wireless Personal Area Networks (LR-
WPANs)", .

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 MAC 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 [I-D.ietf-6lo-ap-nd] should
be more trusted 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.

Future documents may extend this specification by allowing the 6BBR
to redistribute Host routes in routing protocols that would operate
over the Backbone, or in MIPv6 [RFC6275], or FMIP [RFC5568], or the
Locator/ID Separation Protocol (LISP) [RFC6830] to support mobility
on behalf of the 6LNs, etc... LISP may also be used to provide an
equivalent to the EDAR/EDAC exchange using a Map Server / Map
Resolver as a replacement to the 6LBR.

Appendix B. Applicability and Requirements Served

This document specifies proxy-ND functions that can be used to
federate an IPv6 Backbone Link and multiple IPv6 LLNs into a single
Multi-Link Subnet. The proxy-ND functions enable IPv6 ND services
for Duplicate Address Detection (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, covering the types
of networks listed in Appendix B.3 of [RFC8505] "Requirements
Related to Various Low-Power Link Types".

Each LLN in the subnet is attached to an IPv6 Backbone Router
(6BBR). The Backbone Routers interconnect the LLNs and advertise the
Addresses of the 6LNs over the Backbone Link using proxy-ND
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 in Appendix B.1 of [RFC8505] "Requirements Related to Mobility" are met.

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

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 MAC addresses of the 6LNs onto the backbone, keeping the Layer 2 topology simple and stable. This meets the requirements in Appendix B.6 of [RFC8505] "Requirements Related to Scalability", as long has 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 Access Point (AP), or with a Fabric Edge (FE) or a 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 Routing Registrar. The 6LBR can be centralized and either connected to the Backbone Link or reachable over IP. The 6BBR proxy-ND 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 [I-D.nordmark-6man-dad-approaches].

For the TimeSlotted Channel Hopping (TSCH) mode of [IEEEstd802154], the 6TiSCH architecture [I-D.ietf-6tisch-architecture] 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 Appendix B.2 of [RFC8505] "Requirements Related to Routing Protocols".

The registration mechanism may be seen as a more reliable alternate to snooping [I-D.bi-savi-wlan]. It can be noted 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 Multi-Link Subnet [RFC4903] requires avoidance of
multicast/broadcast operations as much as possible even on the
Backbone [I-D.ietf-mboned-ieee802-mcast-problems]. Although hosts
can connect to the Backbone using IPv6 ND operations, multicast RAs
can be saved by using [I-D.ietf-6man-rs-refresh], which also
requires the support of [RFC7559].

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, 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