

6lo
Internet-Draft
Intended status: Standards Track
Expires: April 25, 2019

P. Thubert, Ed.
cisco
C. Perkins
Futurewei
October 22, 2018

IPv6 Backbone Router
draft-ietf-6lo-backbone-router-08

Abstract

Backbone Routers running IPv6 Neighbor Discovery can manage multiple wireless links to form a large MultiLink Subnet, but it is more efficient if IPv6 Neighbor Discovery packets are not broadcast over the wireless links. This specification specifies proxy operations for IPv6 Neighbor Discovery on behalf of devices located on broadcast-inefficient wireless networks. Backbone Routers placed along the wireless edge of the backbone handle IPv6 Neighbor Discovery, and route packets on behalf of registered nodes. Wireless nodes register, or are registered by proxy, to a Backbone Router to establish proxy services in a fashion similar to layer-2 association.

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 April 25, 2019.

Copyright Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of

publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Applicability and Requirements Served	4
3.	Terminology	5
4.	Overview	6
5.	Backbone Router Routing Operations	8
5.1.	Over the Backbone Link	8
5.2.	Proxy Operations Over the LLN Interface	9
5.2.1.	Routing Proxy Operations	10
5.2.2.	Bridging Proxy Operations	10
6.	Backbone Router Proxy Operations	11
6.1.	Primary and Secondary BBRs	12
6.2.	Binding Table	12
6.3.	Registration and Binding Table Entry Creation	13
6.4.	Defending Addresses	14
7.	Security Considerations	15
8.	Protocol Constants	16
9.	IANA Considerations	16
10.	Future Work	16
11.	Acknowledgments	16
12.	References	16
12.1.	Normative References	16
12.2.	Informative References	17
12.3.	External Informative References	19
Appendix A.	Changes from revision 07 to revision 08	20
Authors' Addresses		20

1. Introduction

IEEE STD. 802.1 [[IEEEstd8021](#)] Ethernet Bridging provides an efficient and reliable broadcast service; applications and protocols have been built that heavily depend on that feature for their core operation. Unfortunately, many wireless networks do not economically provide the broadcast capabilities of Ethernet Bridging; protocols designed for bridged networks that rely on broadcast often exhibit disappointing behaviours when applied unmodified to a wireless medium (see [[I-D.ietf-mboned-ieee802-mcast-problems](#)]).

WiFi [[IEEEstd80211](#)] Access Points (APs) deployed in an Extended Service Set (ESS) act as bridges. In order to ensure a solid

connectivity to the devices and protect the medium against harmful broadcasts, they refrain from relying on broadcast-intensive protocols such as Transparent Bridging on the wireless side. Instead, an association process is used to register the MAC addresses of the wireless device (STA) to the AP. The APs subsequently proxy the bridging operation and eliminate the broadcasts.

The IPv6 [[RFC8200](#)] Neighbor Discovery [[RFC4861](#)] [[RFC4862](#)] Protocol (IPv6 ND) operations are reactive and rely heavily on multicast transmissions to locate an on-link correspondent and ensure address uniqueness. Duplicate Address Detection [[RFC4862](#)] (DAD) mechanism was designed as a natural match with the efficient broadcast operation of Ethernet Bridging. However, since broadcast can be unreliable over wireless media, DAD often fails to discover duplications [[I-D.yourtchenko-6man-dad-issues](#)]. DAD usually appears to work on wireless media, not because address duplication is detected and solved as designed, but because the use of 64-bit Interface IDs makes duplication into a very rare event.

IPv6 multicast messages are typically broadcast over the wireless medium. They are processed by most if not all wireless nodes over the ESS fabric even when very few if any of them are subscribed to the multicast address. A simple Neighbor Solicitation (NS) [[RFC4861](#)], that is supposedly targeted to a small group of nodes, can congest the wireless bandwidth [[I-D.ietf-mboned-ieee802-mcast-problems](#)]. The IPv6 ND operation leads to undesirable power consumption in battery-operated devices.

These problems suggest restricting IPv6 ND broadcasts over wireless access links, which can be done by dividing up the subnet. Another way is to take over (proxy) the Layer-3 protocols that rely on broadcast operation at the boundary of the wired and wireless domains, emulating the Layer-2 association at layer-3. For instance, IEEE 802.11 [[IEEEStd80211](#)] specifies ARP and ND proxy [[RFC4389](#)] services at the Access Points (APs).

Current devices rely on snooping for detecting association state, which is failure-prone in lossy and mobile conditions. With snooping, a state (e.g. a new IPv6 address) may not be discovered, or a change of state (e.g. a movement) may be missed, leading to unreliable connectivity.

WPAN devices (i.e., those implementing IEEE STD. 802.15.4 [[IEEEStd802154](#)]) can make use of Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs) [[RFC6775](#)] which treats the wireless medium as different from Ethernet. [RFC 6775](#) is updated as [[I-D.ietf-6lo-rfc6775-update](#)]; the update includes changes that are required by this document.

2. Applicability and Requirements Served

This specification updates and generalizes 6LoWPAN ND to a broader range of Low power and Lossy Networks (LLNs) with support for Duplicate Address Detection (DAD) and address lookup that does not require broadcasts over the LLNs. The term LLN is used loosely in this specification 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, so as to address the requirements listed in [Appendix B.3](#) of [[I-D.ietf-6lo-rfc6775-update](#)] "Requirements Related to the Variety of Low-Power Link types".

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 [[I-D.ietf-6lo-rfc6775-update](#)] "Requirements Related to Routing Protocols".

This specification also applies to wireless links such as Low-Power IEEE STD. 802.11 (Wi-Fi) and IEEE STD. 802.15.1 (Bluetooth) [[IEEEstd802151](#)]. It makes use of extensions to [[RFC6775](#)] to enable proxy operation by the 6BBR, as specified in [[I-D.ietf-6lo-rfc6775-update](#)]. The BBR proxy operations eliminate the need for wireless nodes to respond synchronously when a lookup is performed for their addresses. This provides the function of a Sleep Proxy for ND [[I-D.nordmark-6man-dad-approaches](#)].

This draft establishes a Backbone that treats multiple LLNs as a single IPv6 MultiLink Subnet. Each LLN in the subnet is anchored at an IPv6 Backbone Router (6BBR). The Backbone Routers interconnect the LLNs and advertise the addresses of the 6LNs using proxy-ND operations. This specification extends 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. In this way, mobile nodes may roam rapidly from one 6BBR to the next and requirements in [Appendix B.1](#) of [[I-D.ietf-6lo-rfc6775-update](#)] "Requirements Related to Mobility" are met.

This specification enables any 6LN to register its IPv6 addresses and thereby obtain routing services including proxy-ND operations over the backbone, providing a solution to the requirements expressed in [Appendix B.4](#) of [[I-D.ietf-6lo-rfc6775-update](#)] "Requirements Related to Proxy Operations".

The Link Layer Address (LLA) that is returned as Target LLA (TLA) in Neighbor Advertisements (NA) messages by the 6BBR on behalf of the Registered Node over the backbone may be that of the Registering Node. In that case, the 6BBR needs to bridge the unicast packets (Bridging proxy), or that of the 6BBR on the backbone, in which case the 6BBRs needs to route the unicast packets (Routing proxy). The IPv6 ND operation is minimized as the number of 6LNs grows in the LLN. This meets the requirements in [Appendix B.6](#) of [\[I-D.ietf-6lo-rfc6775-update\]](#) "Requirements Related to Scalability", as long as the 6BBRs are dimensioned for the number of registrations that each needs to support.

In the case of Low-Power IEEE STD. 802.11, a 6BBR may be collocated with a standalone AP or a CAPWAP [\[RFC5415\]](#) wireless controller. Then the wireless client (STA) makes use of this specification to register its IPv6 address(es) to the 6BBR over the wireless medium. In the case RPL, the RPL root is collocated with a 6LoWPAN Border Router (6LBR), and either collocated with or connected to the 6BBR over an IPv6 Link. The 6LBR makes use of this specification to register the 6LNs on their behalf to the 6BBR.

3. Terminology

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](#) [\[RFC2119\]](#) .

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

- o "Neighbor Discovery for IP version 6" [\[RFC4861\]](#),
- o "IPv6 Stateless Address Autoconfiguration" [\[RFC4862\]](#),
- o "Multi-Link Subnet Issues" [\[RFC4903\]](#),
- o "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals" [\[RFC4919\]](#),
- o Neighbor Discovery Optimization for Low-power and Lossy Networks [\[RFC6775\]](#),
- o , "Mobility Support in IPv6" [\[RFC6275\]](#),
- o "Neighbor Discovery Proxies (ND Proxy)" [\[RFC4389\]](#)
- o "Optimistic Duplicate Address Detection" [\[RFC4429\]](#), and

- o "Registration Extensions for 6LoWPAN Neighbor Discovery"
[[I-D.ietf-6lo-rfc6775-update](#)]

This document also uses terminology from [[RFC7102](#)] and [[I-D.ietf-6lo-rfc6775-update](#)], and introduces the following terminology:

Sleeping Proxy

A 6BBR acts as a Sleeping Proxy if it answers ND Neighbor Solicitation over the backbone on behalf of the Registered Node.

Unicasting Proxy

A Unicasting Proxy forwards NS messages to the Registering Node, transforming Layer-2 multicast into unicast.

Routing proxy

A routing proxy advertises its own MAC address as the TLLA in the proxied NAs over the backbone, as opposed to that of the node that performs the registration.

Bridging proxy

A Bridging proxy advertises the MAC address of the node that performs the registration as the TLLA in the proxied NAs over the backbone. In that case, the MAC address and the mobility of 6LN is still visible across the bridged backbone fabric.

Primary BBR

The BBR that will defend a Registered Address for the purpose of DAD over the backbone.

Secondary BBR

A BBR other than the Primary BBR to which an address is registered. A Secondary Router MAY advertise the address over the backbone and proxy for it.

[4.](#) Overview

The services specified in this document assist a 6LN to move freely from an LLN anchored at one 6BBR to an LLN anchored at another 6BBR on the same backbone and keep any or all of the IPv6 addresses that the 6LN has formed.

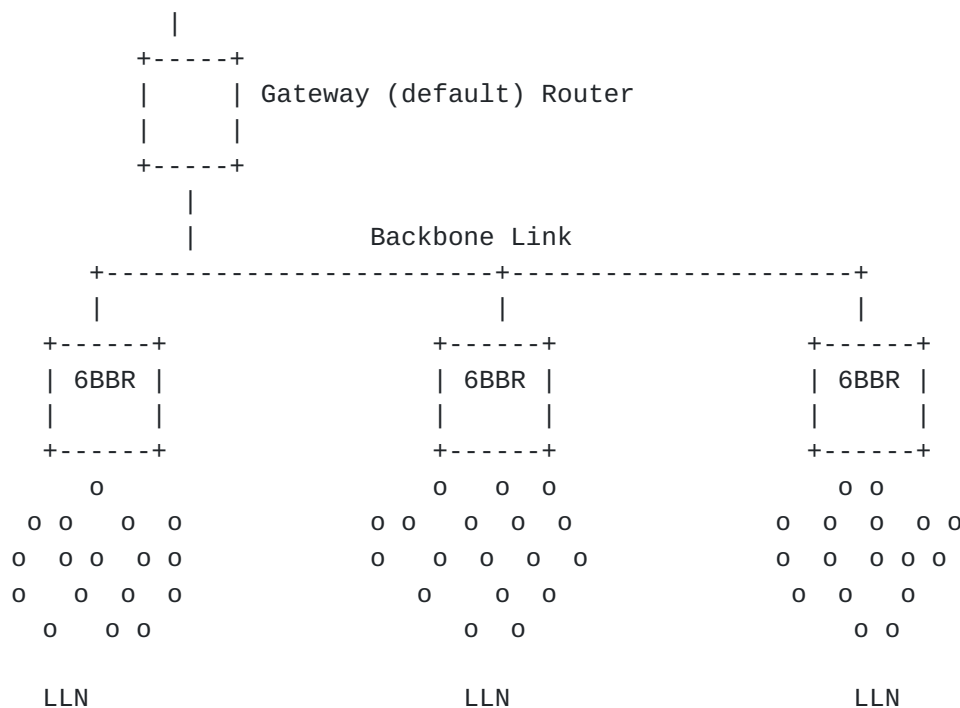


Figure 1: Backbone Link and Backbone Routers

Each Backbone Router (6BBR) maintains a Binding Table of its Registered Nodes. The Binding Tables form a distributed database of wireless 6LNs that reside on the LLNs or on the backbone, and use an extension to IPv6 ND to exchange that information across the Backbone as described below.

The Extended Address Registration Option (EARO) defined in [\[I-D.ietf-6lo-rfc6775-update\]](#) is used in the ND exchanges over the backbone between the 6BBRs to enable the registration for routing and proxy services, as well as distinguish duplication from movement.

Address duplication is detected using the ROVR field in the EARO. In case of conflicting registrations to multiple 6BBRs from the same node, the Transaction ID (TID) in the EARO enables 6BBRs to determine the latest registration for that 6LN.

6BBRs perform ND proxy operations over the backbone, on behalf of their Registered Nodes. Registration to a proxy service is done via a NS/NA(EARO) exchange. 6BBR operation resembles that of a Mobile IPv6 (MIPv6) [\[RFC6275\]](#) Home Agent. This enables mobility support for 6LNs; if they move outside of the network delimited by the Backbone link, then they make use of a Home Agent. Home Agent functionality can easily be collocated with a 6BBR on the same backbone interface of a router.

The Optimistic Duplicate Address Detection [[RFC4429](#)] (ODAD) specification details how an address can be used before a Duplicate Address Detection (DAD) is complete, and mandates that an address that is TENTATIVE should not be associated to a Source Link-Layer Address Option in a Neighbor Solicitation message. This specification makes use of ODAD to create a temporary proxy state in the 6BBR until DAD is completed over the backbone. This way, the specification allows proxy state distribution across multiple 6BBR and co-existence with IPv6 ND over the backbone.

5. Backbone Router Routing Operations

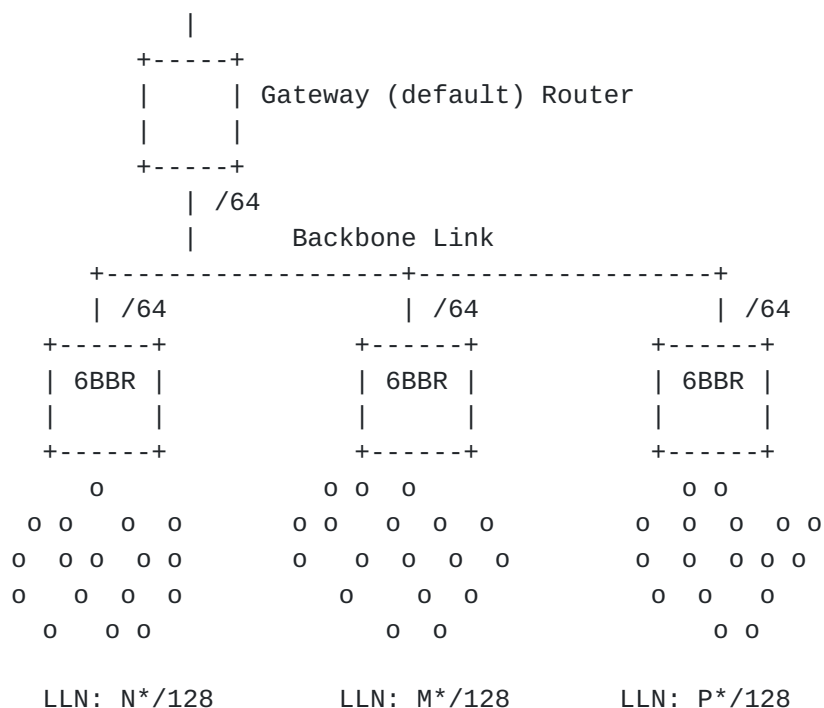


Figure 2: Example Routing Configuration for 3 LLNs in the ML Subnet

5.1. Over the Backbone Link

A 6BBR is a specific kind of Border Router that performs proxy Neighbor Discovery on its backbone interface on behalf of registered 6LLNs on its LLN interfaces.

On the backbone side, the 6BBR advertises the prefixes of the LLNs for which it serves as a proxy. Some restrictions of the attached LLNs will apply to the backbone. In particular, the MTU SHOULD be set to the same value on the backbone and all attached LLNs. The scalability of the multilink subnet [[RFC4903](#)] requires that broadcast operations are avoided as much as possible on the backbone as well.

The 6BBR uses an EARO in the NS-DAD and the multicast NA messages that it generates over the Backbone Link on behalf of a Registered Node. The 6BBR places an EARO in its unicast NA messages, if and only if the NS/NA that stimulates it had an EARO in it and the 'R' bit set.

The 6BBR SHOULD use unicast or the solicited-node multicast address (SNMA) [[RFC4291](#)] to defend its Registered Addresses in its Binding Table over the backbone. In particular, the 6BBR MUST join the SNMA group that corresponds to a Registered Address as soon as it creates an entry for that address, and maintain its SNMA membership as long as it maintains that entry.

Optimistic DAD (ODAD) [[RFC4429](#)] SHOULD be supported by the 6BBRs in their proxy activity over the backbone. A 6BBR supporting ODAD MUST join the SNMA of a Tentative address.

A 6BBR in Routing Proxy mode MAY advertise the Registered IPv6 Address with the 6BBR Link Layer Address, and update Neighbor Cache Entries (NCE) in correspondent nodes over the backbone, using gratuitous NA(Override). This method may fail if the multicast message is not received, and correspondent nodes may maintain an incorrect neighbor state, which they will eventually discover through Neighbor Unreachability Detection (NUD). For slow movements, the NUD procedure defined in [[RFC4861](#)] may time out too quickly, and the support of [[RFC7048](#)] is recommended in all 6LNs in the network.

Multicast should be avoided as much as possible even on the backbone [[I-D.ietf-mboned-ieee802-mcast-problems](#)]. Although hosts can participate using legacy IPv6 ND, all 6LNs connected to the backbone SHOULD support [[I-D.ietf-6man-rs-refresh](#)], which also requires the support of [[RFC7559](#)].

5.2. Proxy Operations Over the LLN Interface

6LNs on the LLN follow [[RFC6775](#)] and do not depend on multicast RAs to discover routers. 6LNs SHOULD accept multicast RAs [[RFC7772](#)], but those are expected to be rare within in the LLN. Nodes SHOULD follow the Simple Procedures for Detecting Network Attachment in IPv6 [[RFC6059](#)] (DNA procedures) to assert movements, and support Packet-Loss Resiliency for Router Solicitations [[RFC7559](#)] to make the unicast RS more reliable.

A 6LN signals that it requires IPv6 ND proxy services from a 6BBR by registering the corresponding IPv6 Address with an NS(EARO) message with the 'R' flag set. The 6LN that performs the registration (the Registering Node) may be the owner of the IPv6 Address (the

Registered Node) or a 6LBR that performs the registration on its behalf.

5.2.1. Routing Proxy Operations

When operating as a Routing Proxy, the BBR installs host routes (/128) to the Registered Addresses within the LLN, via the Registering Node as identified by the Source Address and the SLLA option in the NS(EARO) messages. In that case, the MAC address of the 6LN is not visible at Layer-2 over the backbone. The 6BBR installs a host route towards the Registered Node over the interface toward the 6LN, and routes unicast packets to the 6LN.

The Routing Proxy 6BBR handles the ND protocol over the backbone on behalf of the Registered Nodes, using its own MAC address in the TLLA and SLLA options in proxied NS and NA messages. For each Registered Address, multiple peer Nodes on the backbone may have resolved the address with the 6BBR MAC address, maintaining that mapping in their Neighbor cache.

For each Registered Address, 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 a different 6BBR, the first 6BBR SHOULD unicast a gratuitous NA(Override) to each such peer, to supply the MAC address of the new 6BBR in the TLLA option for the Address.

5.2.2. Bridging Proxy Operations

A Bridging Proxy can be implemented in a Layer-3 switch, or in a wireless Access Point that acts as an IPv6 Host. In the latter case, the SLLA option in the proxied NA messages is that of the Registering Node, and the 6BBR acts as a Layer-2 bridge for unicast packets to the Registered Address. The MAC address in the S/TLLA is that of the Registering Node, which is not necessarily the Registered Node. When a 6LN moves within a LLN mesh, it may attach to a different 6LBR acting as Registering Node, and the MAC address advertised over the backbone might change.

If a registration moves from one 6BBR to the next, but the Registering Node does not change, as indicated by the S/TLLA option in the ND exchanges, there is no need to update the Neighbor Caches of the peer's Nodes on the backbone. On the other hand, if the LLA changes, the 6BBR SHOULD inform all the relevant peers as described above, to update the affected Neighbor Caches. In the same fashion, if the Registering Node changes with a new registration, the 6BBR SHOULD also update the affected Neighbor Caches over the backbone.

6. Backbone Router Proxy Operations

The LLNs attached to each 6BBR are considered different Links in a multi-link subnet. The prefix that is used may still be advertised as on-link on the backbone to support legacy 6LNs. Multicast ND messages are link-scoped and not forwarded across the backbone routers.

By default, a 6BBR operates as a Sleeping Proxy, as follows:

- o Create a new entry in a Binding Table for a new Registered Address and ensure that the address is not a duplicate over the backbone
- o Defend a Registered Address over the backbone using NA messages with the Override bit set on behalf of the sleeping 6LN
- o Advertise a Registered Address over the backbone using NA messages, asynchronously or as a response to a Neighbor Solicitation messages.
- o To deliver packets arriving from the LLN, use Neighbor Solicitation messages to look up the destination over the backbone.
- o Forward packets between the LLN and the backbone.
- o Verify liveness when needed for a stale registration.

A 6BBR may act as a Sleeping Proxy only for a Registered Address that is REACHABLE, or TENTATIVE in which case the answer is delayed. In any other state, the Sleeping Proxy operates as a Unicasting Proxy.

The 6BBR does not act on ND Messages over the backbone unless they are relevant to a Registered Node on the LLN side, saving wireless interference. On the LLN side, the prefixes associated to the MultiLink Subnet are presented as not on-link, so address resolution for other hosts do not occur.

As a Unicasting Proxy, the 6BBR forwards NS lookup messages to the Registering Node, transforming Layer-2 multicast into unicast. This is not possible in UNREACHABLE state, so the NS messages are multicasted, and rate-limited. Retries are possible, using an exponential back-off to protect the medium. In other states, the messages are forwarded to the Registering Node as unicast Layer-2 messages. In TENTATIVE state, the NS message is either held till DAD completes, or dropped if DAD does not complete.

6.1. Primary and Secondary BBRs

A 6BBR MAY be primary or secondary. The primary is the backbone router that has the highest EUI-64 address of all the 6BBRs that share a registration for a 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. The primary Backbone Router is in charge of protecting the address for DAD over the Backbone. Any of the Primary and Secondary 6BBR may claim the address over the backbone, since they are all capable to route from the backbone to the 6LN; the address appears on the backbone as an anycast address.

6.2. Binding Table

Each 6BBR maintains a Binding Table, using IPv6 ND over the backbone to detect duplication. Another document [[I-D.ietf-6lo-rfc6775-update](#)] provides details about how the EARO is used between 6LRs and 6LBRs by way of DAR/DAC messages within the LLN. Addresses in a LLN that can be reachable from the backbone by way of a 6BBR MUST be registered to that 6BBR.

A false positive duplicate detection may arise over the backbone, for instance if a 6LN's Registered Address is registered to more than one LBR, or if the 6LN has moved. Both situations are handled by the 6BBR transparently to the 6LN. In the former case, one LBR becomes primary to defend the address over the backbone while the others become secondary and may still forward packets. In the latter case the LBR that receives the newest registration becomes primary because of the TID.

Only one 6LN may register a given Address at a particular 6BBR. However, that Registered 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 and acknowledged with a status of 4 (TBD); the entry is deleted;

Newer registrations (newer TID, same ROVR, non-null Lifetime) are acknowledged 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 acknowledgement 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 a same Registering Node are acknowledged with a status of 0 (success). If they are not identical, an error SHOULD be logged. In TENTATIVE state, the response is held and may be overwritten, but it MUST be eventually produced and it carries the result of the DAD process;

Older registrations (older TID, same ROVR) from a Registering Node are ignored;

Identical and older registrations (not-newer TID, same ROVR) from a different Registering Node are acknowledged with a status of 3 (moved); this may be rate limited to protect the medium;

Any registration for a different Registered Node (different ROVR) are acknowledged with a status of 1 (duplicate).

6.3. Registration and Binding Table Entry Creation

Upon receiving a registration for a new address with an NS(EARO) with the 'R' bit set, the 6BBR performs DAD over the backbone, placing the new address as target in the NS-DAD message. The EARO from the registration MUST be placed unchanged in the NS-DAD message, and a Neighbor Cache entry created in TENTATIVE state for a duration of TENTATIVE_DURATION. The NS-DAD message is sent multicast over the backbone to the SNMA associated with the registered address, unless that operation is known to be costly, and the 6BBR has an indication from another source (such as a Neighbor Cache entry) that the Registered Address was known on the backbone; in the latter case, an NS-DAD message may be sent as a Layer-2 unicast to the MAC Address that was associated with the Registered Address.

In TENTATIVE state after EARO with 'R' bit set:

1. The entry is removed if an NA is received over the backbone for the Registered Address with no EARO, or containing an EARO with a status of 1 (duplicate) that indicates an existing registration for another 6LN. The ROVR and TID fields in the EARO received over the backbone are ignored. A status of 1 is returned in the EARO of the NA back to the Registering Node;
2. The entry is also removed if an NA with an ARO option with a status of 3 (moved), or a NS with an ARO option that indicates a newer registration for the same Registered Node, is received over

the backbone for the Registered Address. A status of 3 is returned in the NA(EAR0) back to the Registering Node;

3. When a registration is updated but not deleted, e.g. from a newer registration, the DAD process on the backbone continues and the running timers are not restarted;
4. Other NS (including DAD with no EAR0) and NA from the backbone are not acknowledged in TENTATIVE state. To cover legacy 6LNs that do not support ODAD, the list of their origins MAY be stored and then, if the TENTATIVE_DURATION timer elapses, the 6BBR MAY send each such legacy 6LN a unicast NA.
5. When the TENTATIVE_DURATION timer elapses, a status 0 (success) is returned in a NA(EAR0) back to the Registering Node(s), and the entry goes to REACHABLE state for the Registration Lifetime. The 6BBR MUST send a multicast NA(EAR0) to the SNMA associated to the Registered Address over the backbone with the Override bit set so as to take over the binding from other 6BBRs.

6.4. Defending Addresses

If a 6BBR has an entry in REACHABLE state for a Registered Address:

- o If the 6BBR is primary, or does not support the function of primary, it MUST defend that address over the backbone upon receiving NS, either if the NS does not carry an EAR0, or if an EAR0 is present that indicates a different Registering Node (different ROVR). The 6BBR sends a NA message with the Override bit set and the NA carries an EAR0 if and only if the NS-DAD did so. When present, the EAR0 in the NA(Override) that is sent in response to the NS(EAR0) carries a status of 1 (duplicate), and the ROVR and TID fields in the EAR0 are obfuscated with null or random values to avoid network scanning and impersonation attacks.
- o If the 6BBR receives an NS(EAR0) for a newer registration, the 6BBR updates the entry and the routing state to forward packets to the new 6BBR, but keeps the entry REACHABLE. Afterwards, the 6BBR MAY use REDIRECT messages to reroute traffic for the Registered Address to the new 6BBR.
- o If the 6BBR receives an NA(EAR0) for a newer registration, the 6BBR removes its entry and sends a NA(EAR0) with a status of 3 (MOVED) to the Registering Node, if the Registering Node is different from the Registered Node. The 6BBR cleans up existing Neighbor Cache entries in peer nodes as discussed in [Section 5.1](#), by unicasting to each such peer, or one broadcast NA(Override).

- o If the 6BBR receives a NS(LOOKUP) for a Registered Address, it answers immediately with an NA on behalf of the Registered Node, without polling it. There is no need of an EARO in that exchange.
- o When the Registration-Lifetime timer elapses, the entry goes to STALE state for a duration of STABLE_STALE_DURATION in LLNs that keep stable addresses such as LWPANs, and UNSTABLE_STALE_DURATION in LLNs where addresses are renewed rapidly, e.g. for privacy reasons.

The STALE state enables tracking of the backbone peers that have a Neighbor Cache entry 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. In STALE state:

- o If STALE_DURATION elapses, the 6BBR removes the entry.
- o Upon receiving an NA(Override) the 6BBR removes its entry and sends a NA(EARO) with a status of 4 (removed) to the Registering Node.
- o If the 6BBR receives a NS(LOOKUP) for a Registered Address, the 6BBR MUST send an NS(NUD) following rules in [\[RFC7048\]](#) to the Registering Node targeting the Registered Address prior to answering. If the NUD succeeds, the operation in REACHABLE state applies. If the NUD fails, the 6BBR refrains from answering the lookup. The NUD SHOULD be used by the Registering Node to indicate liveness of the Registered Node, if they are different nodes.

7. Security Considerations

This specification applies to LLNS in which the link layer is protected, either by means of physical or IP security for the Backbone Link or MAC sublayer cryptography. In particular, the LLN MAC is required to provide secure unicast to/from the Backbone Router and secure Broadcast from the Backbone Router in a way that prevents tampering with or replaying the RA messages.

The use of EUI-64 for forming the Interface ID in the link local address prevents the usage of Secure ND ([\[RFC3971\]](#) and [\[RFC3972\]](#)) and address privacy techniques. Additional protection against address theft is provided by [\[I-D.ietf-6lo-ap-nd\]](#), which guarantees the ownership of the ROVR.

When the ownership of the ROVR cannot be assessed, this specification limits the cases where the ROVR and the TID are multicasted, and obfuscates them in responses to attempts to take over an address.

8. Protocol Constants

This Specification uses the following constants:

TENTATIVE_DURATION: 800 milliseconds

STABLE_STALE_DURATION: 24 hours

UNSTABLE_STALE_DURATION: 5 minutes

DEFAULT_NS_POLLING: 3 times

9. IANA Considerations

This document has no request to IANA.

10. Future Work

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, or FMIP, or the Locator/ID Separation Protocol (LISP) [[RFC6830](#)] to support mobility on behalf of the 6LNs, etc...

11. Acknowledgments

Kudos to Eric Levy-Abegnoli who designed the First Hop Security infrastructure at Cisco.

12. References

12.1. Normative References

- [I-D.ietf-6lo-rfc6775-update]
Thubert, P., Nordmark, E., Chakrabarti, S., and C. Perkins, "Registration Extensions for 6LoWPAN Neighbor Discovery", [draft-ietf-6lo-rfc6775-update-21](#) (work in progress), June 2018.
- [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>>.
- [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>>.

12.2. Informative References

- [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", [draft-ietf-6lo-ap-nd-08](#) (work in progress), October 2018.
- [I-D.ietf-6man-rs-refresh]
Nordmark, E., Yourtchenko, A., and S. Krishnan, "IPv6 Neighbor Discovery Optional RS/RA Refresh", [draft-ietf-6man-rs-refresh-02](#) (work in progress), October 2016.

[I-D.ietf-6tisch-architecture]

Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", [draft-ietf-6tisch-architecture-15](#) (work in progress), October 2018.

[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", [draft-ietf-mboned-ieee802-mcast-problems-02](#) (work in progress), August 2018.

[I-D.nordmark-6man-dad-approaches]

Nordmark, E., "Possible approaches to make DAD more robust and/or efficient", [draft-nordmark-6man-dad-approaches-02](#) (work in progress), October 2015.

[I-D.yourtchenko-6man-dad-issues]

Yourtchenko, A. and E. Nordmark, "A survey of issues related to IPv6 Duplicate Address Detection", [draft-yourtchenko-6man-dad-issues-01](#) (work in progress), March 2015.

[RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", [RFC 3971](#), DOI 10.17487/RFC3971, March 2005, <<https://www.rfc-editor.org/info/rfc3971>>.

[RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", [RFC 3972](#), DOI 10.17487/RFC3972, March 2005, <<https://www.rfc-editor.org/info/rfc3972>>.

[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>>.

[RFC4919] Kushalnagar, N., Montenegro, G., and C. Schumacher, "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals", [RFC 4919](#), DOI 10.17487/RFC4919, August 2007, <<https://www.rfc-editor.org/info/rfc4919>>.

- [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>>.
- [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>>.
- [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>>.
- [RFC7102] Vasseur, JP., "Terms Used in Routing for Low-Power and Lossy Networks", [RFC 7102](#), DOI 10.17487/RFC7102, January 2014, <<https://www.rfc-editor.org/info/rfc7102>>.
- [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>>.

12.3. External Informative References

- [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. Changes from revision 07 to revision 08

This section lists the changes between [draft-ietf-6lo-backbone-router](#) revisions ...-07.txt and ...-08.txt.

- o Reorganized the order of presentation of some sections so that related material is closer together.
- o Added "Future Work" section.
- o Added this section detailing recent changes.
- o Used '6LN' when LLN node is meant.
- o Updated bibliographic citations.

Authors' Addresses

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

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

Charles E. Perkins
Futurewei
2330 Central Expressway
Santa Clara 95050
United States of America

Email: charliep@computer.org

