

61o
Internet-Draft
Intended status: Standards Track
Expires: May 12, 2016

P. Thubert, Ed.
cisco
November 9, 2015

IPv6 Backbone Router
draft-thubert-61o-backbone-router-03

Abstract

This specification proposes an update to IPv6 Neighbor Discovery, to enhance the operation of IPv6 over wireless links that exhibit lossy multicast support, and enable a large degree of scalability by splitting the broadcast domains. A higher speed backbone federates multiple wireless links to form a large MultiLink Subnet. Backbone Routers acting as Layer-3 Access Point route packets to registered nodes, where an classical Layer-2 Access Point would bridge. Conversely, wireless nodes register to the Backbone Router to setup routing services in a fashion that is essentially similar to a classical 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 <http://datatracker.ietf.org/drafts/current/>.

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

This Internet-Draft will expire on May 12, 2016.

Copyright Notice

Copyright (c) 2015 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 (<http://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 5
- 3. Terminology 6
- 4. Overview 9
- 5. New Types And Formats 10
 - 5.1. Transaction ID 10
 - 5.2. Owner Unique ID 11
 - 5.3. The Enhanced Address Registration Option (EARO) 11
- 6. Backbone Router Routing Operations 14
 - 6.1. Over the Backbone Link 14
 - 6.2. Over the LLN Link 15
- 7. BackBone Router Proxy Operations 17
 - 7.1. Registration and Binding State Creation 19
 - 7.2. Defending Addresses 21
- 8. Security Considerations 22
- 9. Protocol Constants 22
- 10. IANA Considerations 22
- 11. Acknowledgments 23
- 12. References 23
 - 12.1. Normative References 23
 - 12.2. Informative References 24
 - 12.3. External Informative References 28
- Appendix A. Requirements 28
 - A.1. Requirements Related to Mobility 28
 - A.2. Requirements Related to Routing Protocols 29
 - A.3. Requirements Related to the Variety of Low-Power Link types 30
 - A.4. Requirements Related to Proxy Operations 31
 - A.5. Requirements Related to Security 31
 - A.6. Requirements Related to Scalability 32
- Author's Address 33

1. Introduction

Though in most cases, including Low-Power ones, IEEE802.11 [IEEE80211] is operated as a wireless extension to an Ethernet bridged domain, the impact of radio broadcasts for IPv6 [RFC2460] multicast operations, in particular related to the power consumption of battery-operated devices, lead the community to rethink the plain layer-2 approach and consider splitting the broadcast domain between

Thubert

Expires May 12, 2016

[Page 2]

the wired and the wireless access links. To that effect, the current IEEE802.11 specifications require the capability to perform ARP and ND proxy [RFC4389] functions at the Access Points (APs), but rely on snooping for acquiring the related state, which is unsatisfactory in a lossy and mobile environments.

Without a proxy, any IP multicast that circulates in the bridged domain ends up broadcasted by the Access Points to all STAs, including Low-Power battery-operated ones. With an incorrect or missing state in the proxy, a packet may not be delivered to the destination, which may have operational impacts depending on the criticality of the packet.

Some messages are lost for the lack of retries, regardless of their degree of criticality; it results for instance that Duplicate Address Detection (DAD) as defined in [RFC4862] is mostly broken over Wi-Fi [I-D.yourtchenko-6man-dad-issues].

On the other hand, IPv6 multicast messages are processed by most if not all wireless nodes over the fabric even when very few if any of the nodes is effectively listening to the multicast address. It results that a simple Neighbor Solicitation (NS) message [RFC4861], that is supposedly targeted to a very small group of nodes, ends up polluting the whole wireless bandwidth across the fabric [I-D.vyncke-6man-mcast-not-efficient].

It appears that in a variety of Wireless Local Area Networks (WLANs) and Wireless Personal Area Networks (WPANs), the decision to leverage the broadcast support of a particular link should be left to Layer-3 based on the criticality of the message and the number of interested listeners on that link, for the lack of capability to indicate that criticality to the lower layer. To achieve this, the operation at the Access Point cannot be a Layer-2 bridge operation, but that of a Layer-3 router; the concept of MultiLink Subnet (MLSN) must be reintroduced, with IPv6 backbone routers (6BBRs) interconnecting the various links and routing within the subnet. For link-scope multicast operations, a 6BBR participates to MLD on its access links and a multicast routing protocol is setup between the 6BBRs over the backbone of the MLSN.

As the network scales up, none of the approaches of using either broadcast or N*unicast for a multicast packet is really satisfying and the protocols themselves need to be adapted to reduce their use of multicast.

One degree of improvement can be achieved by changing the tuning of the protocol parameters and operational practices, such as suggested in Reducing energy consumption of Router Advertisements

Thubert

Expires May 12, 2016

[Page 3]

[I-D.ietf-v6ops-reducing-ra-energy-consumption] (RA). This works enables to lower the rate of RA messages but does not solve the problem associated with multicast NS and NA messages, which are a lot more frequent in large-scale radio environments with mobile devices which exhibit intermittent access patterns and short-lived IPv6 addresses.

In the context of IEEE802.15.4 [IEEE802154], the more drastic step of considering the radio as a medium that is different from Ethernet because of the impact of multicast, was already taken with the adoption of Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs) [RFC6775]. This specification applies that same thinking to other wireless links such as Low-Power IEEE802.11 (Wi-Fi) and IEEE802.15.1 (Bluetooth) [IEEE802151], and extends [RFC6775] to enable proxy operation by the 6BBR so as to decouple the broadcast domain in the backbone from the wireless links. The proxy operation can be maintained asynchronous so that low-power nodes or nodes that are deep in a mesh do not need to be bothered synchronously when a lookup is performed for their addresses, effectively implementing the ND contribution to the concept of a Sleep Proxy [I-D.nordmark-6man-dad-approaches].

DHCPv6 [RFC3315] is still a viable option in Low power and Lossy Network (LLN) to assign IPv6 global addresses. However, the IETF standard that supports address assignment specifically for LLNs is 6LoWPAN ND [RFC6775], which is a mix of IPv6 stateless autoconfiguration mechanism (SLAAC) [RFC4862] and a new registration process for ND. This specification introduces a Layer-3 association process based on 6LoWPAN ND that maintains a proxy state in the 6BBR to keep the LLN nodes reachable and protect their addresses through sleeping periods.

A number of use cases, including the Industrial Internet, require a large scale deployment of monitoring sensors that can only be realized in a cost-effective fashion with wireless technologies. Mesh networks are deployed when simpler hub-and-spoke topologies are not sufficient for the expected size, throughput, and density. Meshes imply the routing of packets, operated at either Layer-2 or Layer-3. For routing over a mesh at Layer-3, the IETF has designed the IPv6 Routing Protocol over LLN (RPL) [RFC6550]. 6LoWPAN ND was designed as a stand-alone mechanism separately from RPL, and the interaction between the 2 protocols was not defined. This specification details how periodic updates from RPL can be used by the RPL root to renew the association of the RPL node to the 6BBR on its behalf so as to maintain the proxy operation active for that node.

Thubert

Expires May 12, 2016

[Page 4]

This document suggests a limited evolution to [RFC6775] so as to allow operation of a 6LoWPAN ND node while a routing protocol (in first instance RPL) is present and operational. It also suggests a more generalized use of the information in the ARO option of the ND messages outside the strict LLN domain, for instance over a converged backbone.

2. Applicability and Requirements Served

Efficiency aware IPv6 Neighbor Discovery Optimizations [I-D.chakrabarti-nordmark-6man-efficient-nd] suggests that 6LoWPAN ND [RFC6775] can be extended to other types of links beyond IEEE802.15.4 for which it was defined. The registration technique is beneficial when the Link-Layer technique used to carry IPv6 multicast packets is not sufficiently efficient in terms of delivery ratio or energy consumption in the end devices, in particular to enable energy-constrained sleeping nodes. The value of such extension is especially apparent in the case of mobile wireless nodes, to reduce the multicast operations that are related to classical ND ([RFC4861], [RFC4862]) and plague the wireless medium.

This specification updates and generalizes 6LoWPAN ND to a broader range of Low power and Lossy Networks (LLNs) with a solid 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, IEEE802.11AH and IEEE802.15.4 wireless meshes, so as to address the requirements listed in [Appendix A.3](#)

The scope of this draft is a Backbone Link that federates 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 over the Backbone Link and emulate that the LLN nodes are present on the Backbone using proxy-ND operations. This specification extends IPv6 ND over the backbone to discriminate address movement from duplication and eliminate stale state in the backbone routers and backbone nodes once a LLN node has roamed. This way, mobile nodes may roam rapidly from a 6BBR to the next and requirements in [Appendix A.1](#) are met.

This specification can be used by any wireless node to associate at Layer-3 with a 6BBR and register its IPv6 addresses to obtain routing services including proxy-ND operations over the backbone, effectively providing a solution to the requirements expressed in [Appendix A.4](#).

The Link Layer Address (LLA) that is returned as Target LLA (TLLA) in Neighbor Advertisements (NA) messages by the 6BBR on behalf of the

Thubert

Expires May 12, 2016

[Page 5]

Registered Node over the backbone may be that of the Registering Node, in which 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). In the latter case, the 6BBR may maintain the list of correspondents to which it has advertised its own MAC address on behalf of the LLN node and the IPv6 ND operation is minimized as the number of nodes scale up in the LLN. This enables to meet the requirements in [Appendix A.6](#) as long as the 6BBRs are dimensioned for the number of registration that each needs to support.

In the context of the the TimeSlotted Channel Hopping (TSCH) mode of [\[IEEE802154\]](#), the 6TiSCH architecture [\[I-D.ietf-6tisch-architecture\]](#) introduces how a 6LoWPAN ND host could connect to the Internet via a RPL mesh Network, but this requires additions to the 6LoWPAN ND protocol to support mobility and reachability in a secured and manageable environment. This specification details the new operations that are required to implement the 6TiSCH architecture and serves the requirements listed in [Appendix A.2](#).

In the case of Low-Power IEEE802.11, a 6BBR may be collocated with a standalone AP or a CAPWAP [\[RFC5415\]](#) wireless controller, and the wireless client (STA) leverages this specification to register its IPv6 address(es) to the 6BBR over the wireless medium. In the case of a 6TiSCH LLN mesh, 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 leverages this specification to register the LLN nodes on their behalf to the 6BBR. In the case of BTLE, the 6BBR is collocated with the router that implements the BTLE central role as discussed in section 2.2 of [\[I-D.ietf-6lo-btle\]](#).

3. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [\[RFC2119\]](#).

Readers are expected to be familiar with all the terms and concepts that are discussed in "Neighbor Discovery for IP version 6" [\[RFC4861\]](#), "IPv6 Stateless Address Autoconfiguration" [\[RFC4862\]](#), "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals" [\[RFC4919\]](#), "Neighbor Discovery Optimization for Low-power and Lossy Networks" [\[RFC6775\]](#) and "Multi-link Subnet Support in IPv6" [\[I-D.ietf-ipv6-multilink-subnets\]](#).

Readers would benefit from reading "Multi-Link Subnet Issues" [\[RFC4903\]](#), "Mobility Support in IPv6" [\[RFC6275\]](#), "Neighbor Discovery

Thubert

Expires May 12, 2016

[Page 6]

Proxies (ND Proxy)" [[RFC4389](#)] and "Optimistic Duplicate Address Detection" [[RFC4429](#)] prior to this specification for a clear understanding of the art in ND-proxying and binding.

Additionally, this document uses terminology from [[I-D.ietf-roll-terminology](#)] and [[I-D.ietf-6tisch-terminology](#)], and introduces the following terminology:

LLN Low Power Lossy Network. Used loosely in this specification to represent WLANs and WPANs. See [[RFC4919](#)]

Backbone This is an IPv6 transit link that interconnects 2 or more Backbone Routers. It is expected to be deployed as a high speed backbone in order to federate a potentially large set of LLNs. Also referred to as a LLN backbone or Backbone network.

Backbone Router An IPv6 router that federates the LLN using a Backbone link as a backbone. A BBR acts as a 6LoWPAN Border Routers (6LBR) and an Energy Aware Default Router (NEAR).

Extended LLN This is the aggregation of multiple LLNs as defined in [[RFC4919](#)], interconnected by a Backbone Link via Backbone Routers, and forming a single IPv6 MultiLink Subnet.

Registration The process during which a wireless Node registers its address(es) with the Border Router so the 6BBR can proxy ND for it over the backbone.

Binding The state in the 6BBR that associates an IP address with a MAC address, a port and some other information about the node that owns the IP address.

Registered Node The node for which the registration is performed, which owns the fields in the EARO option.

Registering Node The node that performs the registration to the 6BBR, either for one of its own addresses, in which case it is Registered Node and indicates its own MAC Address as SLLA in the NS(ARO), or on behalf of a Registered Node that is reachable over a LLN mesh. In the latter case, if the Registered Node is reachable from the 6BBR over a Mesh-Under mesh, the Registering Node indicates the MAC Address of the Registered Node as SLLA in the NS(ARO). Otherwise, it is expected that the Registered Device is reachable over a Route-Over mesh from the Registering Node, in which case the SLLA in the NS(ARO) is that of the Registering Node, which causes it to attract the packets from the 6BBR to the Registered Node and route them over the LLN.

Registered Address The address owned by the Registered Node node that is being registered.

Sleeping Proxy A 6BBR acts as a Sleeping Proxy if it answers ND Neighbor Solicitation over the backbone on behalf of the Registered Node whenever possible. This is the default mode for this specification but it may be overridden, for instance by configuration, into Unicasting Proxy.

Unicasting Proxy As a Unicasting Proxy, the 6BBR forwards NS messages to the Registering Node, transforming Layer-2 multicast into unicast whenever possible.

Routing proxy A 6BBR acts as a routing proxy if it advertises its own MAC address, as opposed to that of the node that performs the registration, as the TLLA in the proxied NAs over the backbone. In that case, the MAC address of the node is not visible at Layer-2 over the backbone and the bridging fabric is not aware of the addresses of the LLN devices and their mobility. The 6BBR installs a connected host route towards the registered node over the interface to the node, and acts as a Layer-3 router for unicast packets to the node. The 6BBR updates the ND Neighbor Cache Entries (NCE) in correspondent nodes if the wireless node moves and registers to another 6BBR, either with a single broadcast, or with a series of unicast NA(O) messages, indicating the TLLA of the new router.

Bridging proxy A 6BBR acts as a bridging proxy if it 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 the node is still visible across the bridged backbone fabric, as is traditionally the case with Layer-2 APs. The 6BBR acts as a Layer-2 bridge for unicast packets to the registered node. The MAC address exposed in the S/TLLA is that of the Registering Node, which is not necessarily the Registered Device. When a device moves within a LLN mesh, it may end up attached to a different 6LBR acting as Registering Node, and the LLA that is exposed over the backbone will change.

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

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

4. Overview

An LLN node can move freely from an LLN anchored at a Backbone Router to an LLN anchored at another Backbone Router on the same backbone and conserve any of the IPv6 addresses that it has formed, transparently.

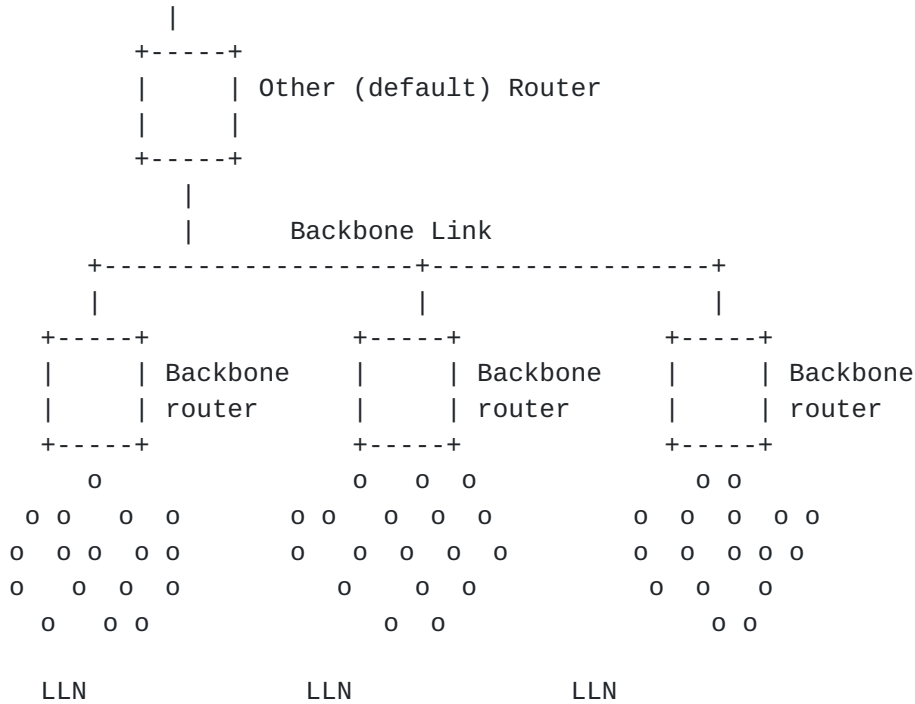


Figure 1: Backbone Link and Backbone Routers

The Backbone Routers maintain an abstract Binding Table of their Registered Nodes. The Binding Table operates as a distributed database of all the wireless Nodes whether they reside on the LLNs or on the backbone, and use an extension to the Neighbor Discovery Protocol to exchange that information across the Backbone in the classical ND reactive fashion.

The Address Registration Option (ARO) defined in [RFC6775] is extended to enable the registration for routing and proxy Neighbor Discovery operations by the 6BBR, and the Extended ARO (EARO) option is included in the ND exchanges over the backbone between the 6BBRs to sort out duplication from movement.

Address duplication is sorted out with the Owner Unique-ID field in the EARO, which is a generalization of the EUI-64 that allows different types of unique IDs beyond the name space derived from the MAC addresses. First-Come First-Serve rules apply, whether the

duplication happens between LLN nodes as represented by their respective 6BBRs, or between an LLN node and a classical node that defends its address over the backbone with classical ND and does not include the EARO option.

In case of conflicting registrations to multiple 6BBRs from a same node, a sequence counter called Transaction ID (TID) is introduced that enables 6BBRs to sort out the latest anchor for that node. Registrations with a same TID are compatible and maintained, but, in case of different TIDs, only the freshest registration is maintained and the stale state is eliminated.

With this specification, Backbone Routers perform ND proxy over the Backbone Link on behalf of their Registered Nodes. The Backbone Router operation is essentially similar to that of a Mobile IPv6 (MIPv6) [RFC6275] Home Agent. This enables mobility support for LLN nodes that would move outside of the network delimited by the Backbone link attach to a Home Agent from that point on. This also enables collocation of Home Agent functionality within Backbone Router functionality on the same backbone interface of a router. Further specification may extend this by allowing the 6BBR to redistribute host routes in routing protocols that would operate over the backbone, or in MIPv6 or the Locator/ID Separation Protocol (LISP) [RFC6830] to support mobility on behalf of the nodes, etc...

The Optimistic Duplicate Address Detection [RFC4429] (ODAD) specification details how an address can be used before a Duplicate Address Detection (DAD) is complete, and insists that an address that is TENTATIVE should not be associated to a Source Link-Layer Address Option in a Neighbor Solicitation message. This specification leverages ODAD to create a temporary proxy state in the 6BBR till DAD is completed over the backbone. This way, the specification enables to distribute proxy states across multiple 6BBR and co-exist with classical ND over the backbone.

5. New Types And Formats

5.1. Transaction ID

The specification expects that the Registered Node can provide a sequence number called Transaction ID (TID) that is incremented with each re-registration. The TID essentially obeys the same rules as the Path Sequence field in the Transit Information Option (TIO) found in RPL's Destination Advertisement Object (DAO). This way, the LLN node can use the same counter for ND and RPL, and a 6LBR acting as RPL root may easily maintain the registration on behalf of a RPL node deep inside the mesh by simply using the RPL TIO Path Sequence as TID for EARO.

Thubert

Expires May 12, 2016

[Page 10]

When a Registered Node is registered to multiple BBRs in parallel, it is expected that the same TID is used, to enable the 6BBRs to correlate the registrations as being a single one, and differentiate that situation from a movement.

If the TIDs are different, the resolution inherited from RPL sorts out the most recent registration and other ones are removed. The operation for computing and comparing the Path Sequence is detailed in [section 7 of \[RFC6550\]](#) and applies to the TID in the exact same fashion.

5.2. Owner Unique ID

The Owner Unique ID (OUID) enables to differentiate a real duplicate address registration from a double registration or a movement. An ND message from the 6BBR over the backbone that is proxied on behalf of a Registered Node must carry the most recent EARO option seen for that node. A NS/NA with an EARO and a NS/NA without a EARO thus represent different nodes and if they relate to a same target then they reflect an address duplication. The Owner Unique ID can be as simple as a EUI-64 burn-in address, if duplicate EUI-64 addresses are avoided.

Alternatively, the unique ID can be a cryptographic string that can be used to prove the ownership of the registration as discussed in Address Protected Neighbor Discovery for Low-power and Lossy Networks [[I-D.sarikaya-6lo-ap-nd](#)].

In any fashion, it is recommended that the node stores the unique Id or the keys used to generate that ID in persistent memory. Otherwise, it will be prevented to re-register after a reboot that would cause a loss of memory until the Backbone Router times out the registration.

5.3. The Enhanced Address Registration Option (EARO)

With the ARO option defined in 6LoWPAN ND [[RFC6775](#)], the address being registered and its owner can be uniquely identified and matched with the Binding Table entries of each Backbone Router.

The Enhanced Address Registration Option (EARO) is intended to be used as a replacement to the ARO option within Neighbor Discovery NS and NA messages between a LLN node and its 6LoWPAN Router (6LR), as well as in Duplicate Address Request (DAR) and the Duplicate Address Confirmation (DAC) messages between 6LRs and 6LBRs in LLNs meshes such as 6TiSCH networks.

An NS message with an EARO option is a registration if and only if it also carries an SLLAO option. The AERO option also used in NS and NA messages between Backbone Routers over the backbone link to sort out the distributed registration state, and in that case, it does not carry the SLLAO option and is not confused with a registration.

The EARO extends the ARO and is recognized by the setting of the TID bit. A node that supports this specification MUST always use an EARO as a replacement to an ARO in its registration to a router. This is harmless since the TID bit and fields are reserved in [RFC6775] are ignored by a legacy router. A router that supports this specification answers to an ARO with an ARO and to an EARO with an EARO.

This specification changes the behavior of the peers in a registration flows. To enable backward compatibility, a node that registers to a router that is not known to support this specification MUST behave as prescribed by [RFC6775]. Once the router is known to support this specification, the node MUST obey this specification.

When using the EARO option, the address being registered is found in the Target Address field of the NS and NA messages. This differs from 6LoWPAN ND [RFC6775] which specifies that the address being registered is the source of the NS.

The reason for this change is to enable proxy-registrations on behalf of other nodes in Route-Over meshes, for instance to enable that a RPL root registers addresses on behalf LLN nodes that are deeper in a 6TiSCH mesh. In that case, the Registering Node MUST indicate its own address as source of the ND message and its MAC address in the Source Link-Layer Address Option (SLLAO), since it still expects to get the packets and route them down the mesh. But the Registered Address belongs to another node, the Registered Node, and that address is indicated in the Target Address field of the NS message.

One way of achieving all the above is for a node to first register an address that it owns in order to validate that the router supports this specification, placing the same address in the Source and Target Address fields of the NS message. The node may for instance register an address that is based on EUI-64. For such address, DAD is not required and using the SLLAO option in the NS is actually more amenable with older ND specifications such as ODAD [RFC4429].

Once that first registration is complete, the node knows from the setting of the TID in the response whether the router supports this specification. If this is verified, the node may register other addresses that it owns, or proxy-register addresses on behalf some another node, indicating those addresses being registered in the

Thubert

Expires May 12, 2016

[Page 12]

Target Address field of the NS messages, while using one of its own, already registered, addresses as source.

The format of the EARO option is as follows:

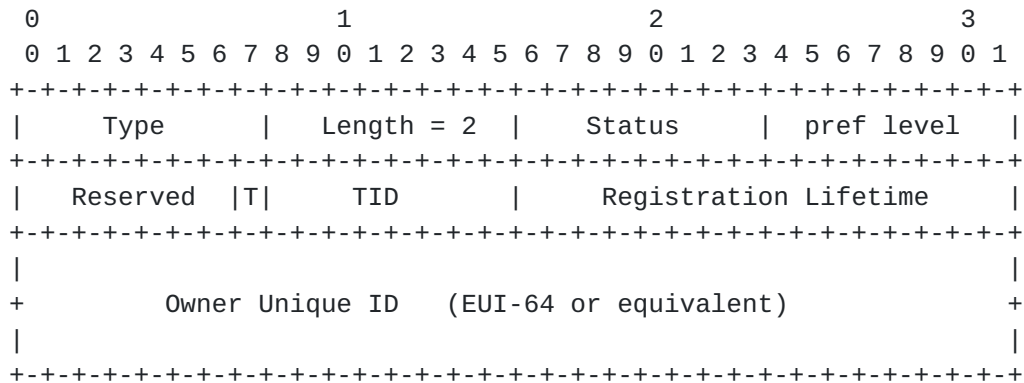


Figure 2: EARO

Option Fields

Type:

Length: 2

Status: OK=0; Duplicate=1; Full=2; Moved=3; Removed=4;

Reserved: This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

T: One bit flag. Set if the next octet is a used as a TID.

TID: 1-byte integer; a transaction id that is maintained by the node and incremented with each transaction. it is recommended that the node maintains the TID in a persistent storage.

Registration Lifetime: 1-byte integer; expressed in minutes. 0 means that the registration has ended and the state should be removed.

Owner Unique Identifier: A globally unique identifier for the node associated. This can be the EUI-64 derived IID of an interface, or some provable ID obtained cryptographically.

6. Backbone Router Routing Operations

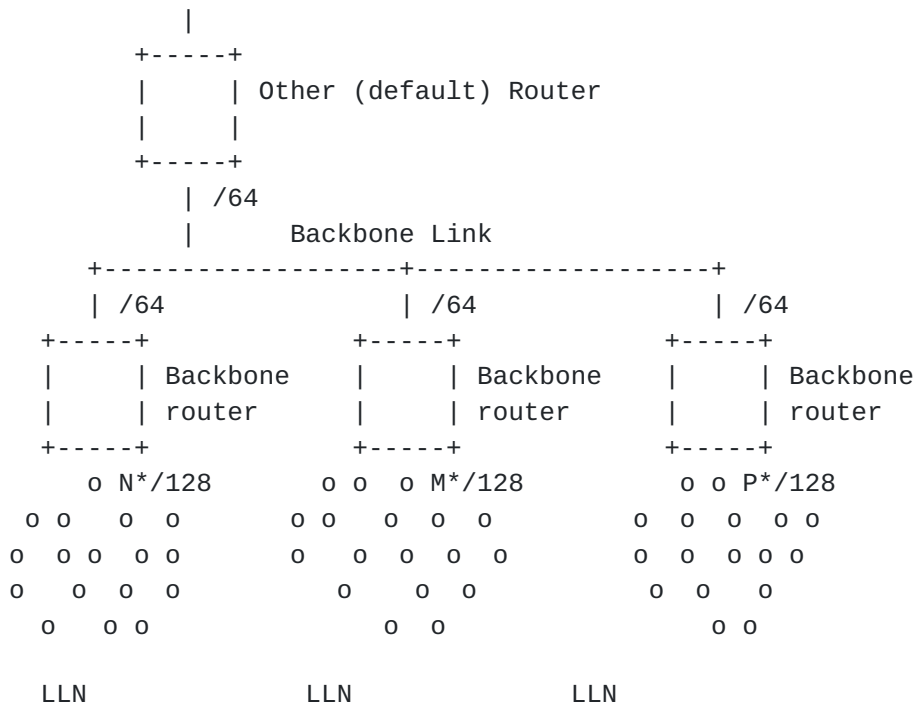


Figure 3: Routing Configuration in the ML Subnet

6.1. Over the Backbone Link

The Backbone Router is a specific kind of Border Router that performs proxy Neighbor Discovery on its backbone interface on behalf of the nodes that it has discovered on its LLN interfaces.

The backbone is expected to be a high speed, reliable Backbone link, with affordable and reliable multicast capabilities, such as a bridged Ethernet Network, and to allow a full support of classical ND as specified in [RFC4861] and subsequent RFCs. In other words, the backbone is not a LLN.

Still, some restrictions of the attached LLNs will apply to the backbone. In particular, it is expected that the MTU is set to the same value on the backbone and all attached LLNs, and the scalability of the whole subnet requires that broadcast operations are avoided as much as possible on the backbone as well. Unless configured otherwise, the Backbone Router MUST echo the MTU that it learns in RAs over the backbone in the RAs that it sends towards the LLN links.

As a router, the Backbone Router behaves like any other IPv6 router on the backbone side. It has a connected route installed towards the

backbone for the prefixes that are present on that backbone and that it proxies for on the LLN interfaces.

As a proxy, the 6BBR uses an EARO option in the NS-DAD and the multicast NA messages that it generates on behalf of a Registered Node, and it places an EARO in its unicast NA messages if and only if the NS/NA that stimulates it had an EARO in it.

When possible, the 6BBR SHOULD use unicast or solicited-node multicast address (SNMA) [RFC4291] to defend its Registered Addresses 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 as long as it maintains that entry, whatever the state of the entry. The expectation is that it is possible to get a message delivered to all the nodes on the backbone that listen to a particular address and support this specification - which includes all the 6BBRs in the MultiLink Subnet - by sending a multicast message to the associated SNMA over the backbone.

The support of Optimistic DAD (ODAD) [RFC4429] is recommended for all nodes in the backbone and followed by the 6BBRs in their proxy activity over the backbone. With ODAD, any optimistic node MUST join the SNMA of a Tentative address, which interacts better with this specification.

This specification allows the 6BBR in Routing Proxy mode to advertise the Registered IPv6 Address with the 6BBR Link Layer Address, and attempts to 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 properly received, and correspondent nodes may maintain an incorrect neighbor state, which they will eventually discover through Neighbor Unreachability Detection (NUD). Because mobility may be slow, the NUD procedure defined in [RFC4861] may be too impatient, and the support of [RFC7048] is recommended in all nodes in the network.

Since the MultiLink Subnet may grow very large in terms of individual IPv6 addresses, multicasts should be avoided as much as possible even on the backbone. Though it is possible for plain hosts to participate with legacy IPv6 ND support, the support by all nodes connected to the backbone of [I-D.nordmark-6man-rs-refresh] is recommended, and this implies the support of [RFC7559] as well.

6.2. Over the LLN Link

As a router, the Nodes and Backbone Router operation on the LLN follows [RFC6775]. Per that specification, LLN Hosts generally do not depend on multicast RAs to discover routers. It is still

generally required for LLN nodes to accept multicast RAs [[I-D.ietf-v6ops-reducing-ra-energy-consumption](#)], but those are rare on the LLN link. Nodes are expected to follow the Simple Procedures for Detecting Network Attachment in IPv6 [[RFC6059](#)] (DNA procedures) to assert movements, and to support the Packet-Loss Resiliency for Router Solicitations [[RFC7559](#)] to make the unicast RS more reliable.

The Backbone Router acquires its states about the addresses on the LLN side through a registration process from either the nodes themselves, or from a node such as a RPL root / 6LBR (the Registering Node) that performs the registration on behalf of the address owner (the Registered Node).

When operating as a Routing Proxy, the router installs hosts routes (/128) to the Registered Addresses over the LLN links, via the Registering Node as identified by the Source Address and the SLLAO option in the NS(EARO) messages.

In that mode, the 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. It results that for each Registered Address, a number of peer Nodes on the backbone have resolved the address with the 6BBR MAC address and keep that mapping stored in their Neighbor cache.

The 6BBR SHOULD maintain, per Registered Address, the list of the peers on the backbone to which it answered with its MAC address, and when a binding moves to a different 6BBR, it SHOULD send a unicast gratuitous NA(0) individually to each of them to inform them that the address has moved and pass the MAC address of the new 6BBR in the TLLAO option. If the 6BBR can not maintain that list, then it SHOULD remember whether that list is empty or not and if not, send a multicast NA(0) to all nodes to update the impacted Neighbor Caches with the information from the new 6BBR.

The Bridging Proxy is a variation where the BBR function is implemented in a Layer-3 switch or an wireless Access Point that acts as a Host from the IPv6 standpoint, and, in particular, does not operate the routing of IPv6 packets. In that case, the SLLAO in the proxied NA messages is that of the Registering Node and classical bridging operations take place on data frames.

If a registration moves from one 6BBR to the next, but the Registering Node does not change, as indicated by the S/TLLAO option in the ND exchanges, there is no need to update the Neighbor Caches in the peers Nodes on the backbone. On the other hand, if the LLAO changes, the 6BBR SHOULD inform all the relevant peers as described above, to update the impacted Neighbor Caches. In the same fashion,

if the Registering Node changes with a new registration, the 6BBR SHOULD also update the impacted Neighbor Caches over the backbone.

7. Backbone Router Proxy Operations

This specification enables a Backbone Router to proxy Neighbor Discovery operations over the backbone on behalf of the nodes that are registered to it, allowing any node on the backbone to reach a Registered Node as if it was on-link. The backbone and the LLNs are considered different Links in a MultiLink subnet but the prefix that is used may still be advertised as on-link on the backbone to support legacy nodes; multicast ND messages are link-scoped and not forwarded across the backbone routers.

ND Messages on the backbone side that do not match to a registration on the LLN side are not acted upon on the LLN side, which stands protected. 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.

The default operation in this specification is Sleeping proxy which means:

- o creating a new entry in an abstract Binding Table for a new Registered Address and validating that the address is not a duplicate over the backbone
- o defending a Registered Address over the backbone using NA messages with the Override bit set on behalf of the sleeping node whenever possible
- o advertising a Registered Address over the backbone using NA messages, asynchronously or as a response to a Neighbor Solicitation messages.
- o Looking up a destination over the backbone in order to deliver packets arriving from the LLN using Neighbor Solicitation messages.
- o Forwarding packets from the LLN over the backbone, and the other way around.
- o Eventually triggering a liveness verification of a stale registration.

A 6BBR may act as a Sleeping Proxy only if the state of the binding entry is REACHABLE, or TENTATIVE in which case the answer is delayed.

In any other state, the Sleeping Proxy operates as a Unicasting Proxy.

As a Unicasting Proxy, the 6BBR forwards NS messages to the Registering Node, transforming Layer-2 multicast into unicast whenever possible. This is not possible in UNREACHABLE state, so the NS messages are multicasted, and rate-limited to protect the medium with an exponential back-off. 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.

The draft introduces the optional concept of primary and secondary BBRs. 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 Owner Unique ID and same Transaction ID, the EUI-64 address being considered as an unsigned 64bit integer. The concept is defined with the granularity of an address, that is a given 6BBR can be primary for a given address and secondary or another one, regardless on whether the addresses belong to the same node or not. 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 LLN node, and the address appears on the backbone as an anycast address.

The Backbone Routers maintain a distributed binding table, using classical ND over the backbone to detect duplication. This specification requires that:

1. All addresses that can be reachable from the backbone, including IPv6 addresses based on burn-in EUI64 addresses MUST be registered to the 6BBR.
2. A Registered Node MUST include the EAR0 option in an NS message that used to register an addresses to a 6LR; the 6LR MUST propagate that option unchanged to the 6LBR in the DAR/DAC exchange, and the 6LBR MUST propagate that option unchanged in proxy registrations.
3. The 6LR MUST echo the same EAR0 option in the NA that it uses to respond, but for the status filed which is not used in NS messages, and significant in NA.

A false positive duplicate detection may arise over the backbone, for instance if the Registered Address is registered to more than one LBR, or if the node has moved. Both situations are handled gracefully unbeknownst to the node. 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 back and forth. In the latter case the LBR that receives the newest registration wins and becomes primary.

The expectation in this specification is that there is a single Registering Node at a time per Backbone Router for a given Registered Address, but that a Registered Address may be registered to Multiple 6BBRs for higher availability.

Over the LLN, and for any given Registered Address, it is REQUIRED that:

de-registrations (newer TID, same OUID, null Lifetime) are accepted and responded immediately with a status of 4; the entry is deleted;

newer registrations (newer TID, same OUID, non-null Lifetime) are accepted and responded with a status of 0 (success); the entry is updated with the new TID, the new Registration Lifetime and the new Registering Node, if any has changed; in TENTATIVE state the 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 OUID) from a same Registering Node are not processed but responded with a status of 0 (success); they are expected to be identical and an error may be logged if not; 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 (not(newer or equal) TID, same OUID) from a same Registering Node are ignored;

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

and any registration for a different Registered Node (different OUID) are responded immediately with a status of 1 (duplicate).

7.1. Registration and Binding State Creation

Upon a registration for a new address with an NS(EARO), the 6BBR performs a DAD operation 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 an entry is created in TENTATIVE state for a duration of TENTATIVE_DURATION. The NS-DAD message is sent multicast over the backbone to the SNMA address associated with the registered address. If that operation is known to be costly, and the 6BBR has an indication from another source (such as a NCE) that the Registered Address was present on the backbone, that information may be leveraged to send the NS-DAD message as a Layer-2 unicast to the MAC that was associated with the Registered Address.

In TENTATIVE state:

- o the entry is removed if an NA is received over the backbone for the Registered Address with no EARO option, or with an EARO option with a status of 1 (duplicate) that indicates an existing registration for another LLN node. The OUID and TID fields in the EARO option received over the backbone are ignored. A status of 1 is returned in the EARO option of the NA back to the Registering Node;
- o the entry is also removed if an NA with an ARO option with a status of 3 (moved), or a NS-DAD 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(EARO) back to the Registering Node;
- o 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;
- o Other NS (including DAD with no EARO option) and NA from the backbone are not responded in TENTATIVE state, but the list of their origins may be kept in memory and if so, the 6BBR may send them each a unicast NA with eventually an EARO option when the TENTATIVE_DURATION timer elapses, so as to cover legacy nodes that do not support ODAD.
- o When the TENTATIVE_DURATION timer elapses, a status 0 (success) is returned in a NA(EARO) back to the Registering Node(s), and the entry goes to REACHABLE state for the Registration Lifetime; the DAD process is successful and the 6BBR MUST send a multicast NA(EARO) 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.

7.2. 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 concept, it MUST defend that address over the backbone upon an incoming NS-DAD, either if the NS does not carry an EARO, or if an EARO is present that indicates a different Registering Node (different OUID). The 6BBR sends a NA message with the Override bit set and the NA carries an EARO option if and only if the NS-DAD did so. When present, the EARO in the NA(0) that is sent in response to the NS-DAD(EARO) carries a status of 1 (duplicate), and the OUID and TID fields in the EARO option are obfuscated with null or random values to avoid network scanning and impersonation attacks.
- o If the 6BBR receives an NS-DAD(EARO) that reflect a newer registration, the 6BBR updates the entry and the routing state to forward packets to the new 6BBR, but keeps the entry REACHABLE. In that phase, it MAY use REDIRECT messages to reroute traffic for the Registered Address to the new 6BBR.
- o If the 6BBR receives an NA(EARO) that reflect a newer registration, the 6BBR removes its entry and sends a NA(AERO) with a status of 3 (moved) to the Registering Node, if the Registering Node is different from the Registered Node. If necessary, the 6BBR cleans up ND cache in peers nodes as discussed in [Section 6.1](#), by sending a series of unicast to the impacted nodes, or one broadcast NA(0) to all-nodes.
- o If the 6BBR received 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 is a chance to keep track of the backbone peers that may have an ND cache pointing on this 6BBR in case the Registered Address shows back up on this or a different 6BBR at a later time. In STALE state:

- o If the Registered Address is claimed by another node on the backbone, with an NS-DAD or an NA, the 6BBR does not defend the address. Upon an NA(0), or the stale time elapses, the 6BBR

removes its entry and sends a NA(AERO) with a status of 4 (removed) to the Registering Node.

- o If the 6BBR received 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 expected to be mapped by the Registering Node into a liveness validation of the Registered Node if they are in fact different nodes.

8. Security Considerations

This specification expects that the link layer is sufficiently protected, either by means of physical or IP security for the Backbone Link or MAC sublayer cryptography. In particular, it is expected that the LLN MAC provides secure unicast to/from the Backbone Router and secure Broadcast from the Backbone Router in a way that prevents tempering 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. This specification RECOMMENDS the use of additional protection against address theft such as provided by [I-D.sarikaya-6lo-ap-nd], which guarantees the ownership of the OUID.

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

9. 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

10. IANA Considerations

This document requires the following additions:

Address Registration Option Status Values Registry

Status	Description
3	Moved: The registration fails because it is not the freshest.
4	Removed: The binding state was removed

IANA is required to change the registry accordingly

Table 1: New ARO Status values

11. Acknowledgments

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

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), DOI 10.17487/RFC2460, December 1998, <<http://www.rfc-editor.org/info/rfc2460>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 4291](#), DOI 10.17487/RFC4291, February 2006, <<http://www.rfc-editor.org/info/rfc4291>>.
- [RFC4429] Moore, N., "Optimistic Duplicate Address Detection (DAD) for IPv6", [RFC 4429](#), DOI 10.17487/RFC4429, April 2006, <<http://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, <<http://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, <<http://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, <<http://www.rfc-editor.org/info/rfc6059>>.
- [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, <<http://www.rfc-editor.org/info/rfc6550>>.
- [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, <<http://www.rfc-editor.org/info/rfc6775>>.

12.2. Informative References

- [I-D.chakrabarti-nordmark-6man-efficient-nd]
Chakrabarti, S., Nordmark, E., Thubert, P., and M. Wasserman, "IPv6 Neighbor Discovery Optimizations for Wired and Wireless Networks", [draft-chakrabarti-nordmark-6man-efficient-nd-07](#) (work in progress), February 2015.
- [I-D.delcarpio-6lo-wlanah]
Vega, L., Robles, I., and R. Morabito, "IPv6 over 802.11ah", [draft-delcarpio-6lo-wlanah-01](#) (work in progress), October 2015.
- [I-D.ietf-6lo-6lobac]
Lynn, K., Martocci, J., Neilson, C., and S. Donaldson, "Transmission of IPv6 over MS/TP Networks", [draft-ietf-6lo-6lobac-03](#) (work in progress), October 2015.
- [I-D.ietf-6lo-btle]
Nieminen, J., Savolainen, T., Isomaki, M., Patil, B., Shelby, Z., and C. Gomez, "IPv6 over BLUETOOTH(R) Low Energy", [draft-ietf-6lo-btle-17](#) (work in progress), August 2015.

[I-D.ietf-6lo-dect-ule]

Mariager, P., Petersen, J., Shelby, Z., Logt, M., and D. Barthel, "Transmission of IPv6 Packets over DECT Ultra Low Energy", [draft-ietf-6lo-dect-ule-03](#) (work in progress), September 2015.

[I-D.ietf-6lo-nfc]

Youn, J. and Y. Hong, "Transmission of IPv6 Packets over Near Field Communication", [draft-ietf-6lo-nfc-02](#) (work in progress), October 2015.

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

Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", [draft-ietf-6tisch-architecture-08](#) (work in progress), May 2015.

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

Palattella, M., Thubert, P., Watteyne, T., and Q. Wang, "Terminology in IPv6 over the TSCH mode of IEEE 802.15.4e", [draft-ietf-6tisch-terminology-06](#) (work in progress), November 2015.

[I-D.ietf-bier-architecture]

Wijnands, I., Rosen, E., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast using Bit Index Explicit Replication", [draft-ietf-bier-architecture-02](#) (work in progress), July 2015.

[I-D.ietf-ipv6-multilink-subnets]

Thaler, D. and C. Huitema, "Multi-link Subnet Support in IPv6", [draft-ietf-ipv6-multilink-subnets-00](#) (work in progress), July 2002.

[I-D.ietf-roll-terminology]

Vasseur, J., "Terms used in Routing for Low power And Lossy Networks", [draft-ietf-roll-terminology-13](#) (work in progress), October 2013.

[I-D.ietf-v6ops-reducing-ra-energy-consumption]

Yourtchenko, A. and L. Colitti, "Reducing energy consumption of Router Advertisements", [draft-ietf-v6ops-reducing-ra-energy-consumption-03](#) (work in progress), November 2015.

[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.nordmark-6man-rs-refresh]

Nordmark, E., Yourtchenko, A., and S. Krishnan, "IPv6 Neighbor Discovery Optional Unicast RS/RA Refresh", [draft-nordmark-6man-rs-refresh-01](#) (work in progress), October 2014.

[I-D.popa-6lo-6loplcv6-over-ieee19012-networks]

Popa, D. and J. Hui, "6LoPLC: Transmission of IPv6 Packets over IEEE 1901.2 Narrowband Powerline Communication Networks", [draft-popa-6lo-6loplcv6-over-ieee19012-networks-00](#) (work in progress), March 2014.

[I-D.sarikaya-6lo-ap-nd]

Sarikaya, B. and P. Thubert, "Address Protected Neighbor Discovery for Low-power and Lossy Networks", [draft-sarikaya-6lo-ap-nd-01](#) (work in progress), October 2015.

[I-D.vyncke-6man-mcast-not-efficient]

Vyncke, E., Thubert, P., Levy-Abegnoli, E., and A. Yourtchenko, "Why Network-Layer Multicast is Not Always Efficient At Datalink Layer", [draft-vyncke-6man-mcast-not-efficient-01](#) (work in progress), February 2014.

[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.

[RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3315](#), DOI 10.17487/RFC3315, July 2003, <<http://www.rfc-editor.org/info/rfc3315>>.

[RFC3810] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", [RFC 3810](#), DOI 10.17487/RFC3810, June 2004, <<http://www.rfc-editor.org/info/rfc3810>>.

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

[RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", [RFC 3972](#), DOI 10.17487/RFC3972, March 2005, <<http://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, <<http://www.rfc-editor.org/info/rfc4389>>.
- [RFC4903] Thaler, D., "Multi-Link Subnet Issues", [RFC 4903](#), DOI 10.17487/RFC4903, June 2007, <<http://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, <<http://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, <<http://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, <<http://www.rfc-editor.org/info/rfc6275>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", [RFC 6282](#), DOI 10.17487/RFC6282, September 2011, <<http://www.rfc-editor.org/info/rfc6282>>.
- [RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", [RFC 6830](#), DOI 10.17487/RFC6830, January 2013, <<http://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, <<http://www.rfc-editor.org/info/rfc7048>>.
- [RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", [RFC 7217](#), DOI 10.17487/RFC7217, April 2014, <<http://www.rfc-editor.org/info/rfc7217>>.

- [RFC7428] Brandt, A. and J. Buron, "Transmission of IPv6 Packets over ITU-T G.9959 Networks", [RFC 7428](#), DOI 10.17487/RFC7428, February 2015, <<http://www.rfc-editor.org/info/rfc7428>>.
- [RFC7559] Krishnan, S., Anipko, D., and D. Thaler, "Packet-Loss Resiliency for Router Solicitations", [RFC 7559](#), DOI 10.17487/RFC7559, May 2015, <<http://www.rfc-editor.org/info/rfc7559>>.

12.3. External Informative References

- [IEEE80211] 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".
- [IEEE802151] 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)".
- [IEEE802154] 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. Requirements

This section lists requirements that were discussed at 6lo for an update to 6LoWPAN ND. This specification meets most of them, but those listed in [Appendix A.5](#) which are deferred to a different specification such as [[I-D.sarikaya-6lo-ap-nd](#)].

A.1. Requirements Related to Mobility

Due to the unstable nature of LLN links, even in a LLN of immobile nodes a 6LoWPAN Node may change its point of attachment to a 6LR, say 6LR-a, and may not be able to notify 6LR-a. Consequently, 6LR-a may still attract traffic that it cannot deliver any more. When links to

a 6LR change state, there is thus a need to identify stale states in a 6LR and restore reachability in a timely fashion.

Req1.1: Upon a change of point of attachment, connectivity via a new 6LR MUST be restored timely without the need to de-register from the previous 6LR.

Req1.2: For that purpose, the protocol MUST enable to differentiate between multiple registrations from one 6LoWPAN Node and registrations from different 6LoWPAN Nodes claiming the same address.

Req1.3: Stale states MUST be cleaned up in 6LRs.

Req1.4: A 6LoWPAN Node SHOULD also be capable to register its Address to multiple 6LRs, and this, concurrently.

A.2. Requirements Related to Routing Protocols

The point of attachment of a 6LoWPAN Node may be a 6LR in an LLN mesh. IPv6 routing in a LLN can be based on RPL, which is the routing protocol that was defined at the IETF for this particular purpose. Other routing protocols than RPL are also considered by Standard Defining Organizations (SDO) on the basis of the expected network characteristics. It is required that a 6LoWPAN Node attached via ND to a 6LR would need to participate in the selected routing protocol to obtain reachability via the 6LR.

Next to the 6LBR unicast address registered by ND, other addresses including multicast addresses are needed as well. For example a routing protocol often uses a multicast address to register changes to established paths. ND needs to register such a multicast address to enable routing concurrently with discovery.

Multicast is needed for groups. Groups MAY be formed by device type (e.g. routers, street lamps), location (Geography, RPL sub-tree), or both.

The Bit Index Explicit Replication (BIER) Architecture [[I-D.ietf-bier-architecture](#)] proposes an optimized technique to enable multicast in a LLN with a very limited requirement for routing state in the nodes.

Related requirements are:

Req2.1: The ND registration method SHOULD be extended in such a fashion that the 6LR MAY advertise the Address of a 6LoWPAN Node over the selected routing protocol and obtain reachability to that Address using the selected routing protocol.

Req2.2: Considering RPL, the Address Registration Option that is used in the ND registration SHOULD be extended to carry enough information to generate a DAO message as specified in [\[RFC6550\]](#) section 6.4, in particular the capability to compute a Path Sequence and, as an option, a RPLInstanceID.

Req2.3: Multicast operations SHOULD be supported and optimized, for instance using BIER or MPL. Whether ND is appropriate for the registration to the 6BBR is to be defined, considering the additional burden of supporting the Multicast Listener Discovery Version 2 [\[RFC3810\]](#) (MLDv2) for IPv6.

A.3. Requirements Related to the Variety of Low-Power Link types

6LoWPAN ND [\[RFC6775\]](#) was defined with a focus on IEEE802.15.4 and in particular the capability to derive a unique Identifier from a globally unique MAC-64 address. At this point, the 6lo Working Group is extending the 6LoWPAN Header Compression (HC) [\[RFC6282\]](#) technique to other link types ITU-T G.9959 [\[RFC7428\]](#), Master-Slave/Token-Passing [\[I-D.ietf-6lo-6lobac\]](#), DECT Ultra Low Energy [\[I-D.ietf-6lo-dect-ule\]](#), Near Field Communication [\[I-D.ietf-6lo-nfc\]](#), IEEE802.11ah [\[I-D.delcarpio-6lo-wlanah\]](#), as well as IEEE1901.2 Narrowband Powerline Communication Networks [\[I-D.popa-6lo-6loplc-ipv6-over-ieee19012-networks\]](#) and BLUETOOTH(R) Low Energy [\[I-D.ietf-6lo-btle\]](#).

Related requirements are:

Req3.1: The support of the registration mechanism SHOULD be extended to more LLN links than IEEE 802.15.4, matching at least the LLN links for which an "IPv6 over foo" specification exists, as well as Low-Power Wi-Fi.

Req3.2: As part of this extension, a mechanism to compute a unique Identifier should be provided, with the capability to form a Link-Local Address that SHOULD be unique at least within the LLN connected to a 6LBR discovered by ND in each node within the LLN.

Req3.3: The Address Registration Option used in the ND registration SHOULD be extended to carry the relevant forms of unique Identifier.

Req3.4: The Neighbour Discovery should specify the formation of a site-local address that follows the security recommendations from [\[RFC7217\]](#).

A.4. Requirements Related to Proxy Operations

Duty-cycled devices may not be able to answer themselves to a lookup from a node that uses classical ND on a backbone and may need a proxy. Additionally, the duty-cycled device may need to rely on the 6LBR to perform registration to the 6BBR.

The ND registration method SHOULD defend the addresses of duty-cycled devices that are sleeping most of the time and not capable to defend their own Addresses.

Related requirements are:

Req4.1: The registration mechanism SHOULD enable a third party to proxy register an Address on behalf of a 6LoWPAN node that may be sleeping or located deeper in an LLN mesh.

Req4.2: The registration mechanism SHOULD be applicable to a duty-cycled device regardless of the link type, and enable a 6BBR to operate as a proxy to defend the registered Addresses on its behalf.

Req4.3: The registration mechanism SHOULD enable long sleep durations, in the order of multiple days to a month.

A.5. Requirements Related to Security

In order to guarantee the operations of the 6LoWPAN ND flows, the spoofing of the 6LR, 6LBR and 6BBRs roles should be avoided. Once a node successfully registers an address, 6LoWPAN ND should provide energy-efficient means for the 6LBR to protect that ownership even when the node that registered the address is sleeping.

In particular, the 6LR and the 6LBR then should be able to verify whether a subsequent registration for a given Address comes from the original node.

In a LLN it makes sense to base security on layer-2 security. During bootstrap of the LLN, nodes join the network after authorization by a Joining Assistant (JA) or a Commissioning Tool (CT). After joining nodes communicate with each other via secured links. The keys for the layer-2 security are distributed by the JA/CT. The JA/CT can be part of the LLN or be outside the LLN. In both cases it is needed that packets are routed between JA/CT and the joining node.

Related requirements are:

Req5.1: 6LoWPAN ND security mechanisms SHOULD provide a mechanism for the 6LR, 6LBR and 6BBR to authenticate and authorize one another for

their respective roles, as well as with the 6LoWPAN Node for the role of 6LR.

Req5.2: 6LoWPAN ND security mechanisms SHOULD provide a mechanism for the 6LR and the 6LBR to validate new registration of authorized nodes. Joining of unauthorized nodes MUST be impossible.

Req5.3: 6LoWPAN ND security mechanisms SHOULD lead to small packet sizes. In particular, the NS, NA, DAR and DAC messages for a re-registration flow SHOULD NOT exceed 80 octets so as to fit in a secured IEEE802.15.4 frame.

Req5.4: Recurrent 6LoWPAN ND security operations MUST NOT be computationally intensive on the LoWPAN Node CPU. When a Key hash calculation is employed, a mechanism lighter than SHA-1 SHOULD be preferred.

Req5.5: The number of Keys that the 6LoWPAN Node needs to manipulate SHOULD be minimized.

Req5.6: The 6LoWPAN ND security mechanisms SHOULD enable CCM* for use at both Layer 2 and Layer 3, and SHOULD enable the reuse of security code that has to be present on the device for upper layer security such as TLS.

Req5.7: Public key and signature sizes SHOULD be minimized while maintaining adequate confidentiality and data origin authentication for multiple types of applications with various degrees of criticality.

Req5.8: Routing of packets should continue when links pass from the unsecured to the secured state.

Req5.9: 6LoWPAN ND security mechanisms SHOULD provide a mechanism for the 6LR and the 6LBR to validate whether a new registration for a given address corresponds to the same 6LoWPAN Node that registered it initially, and, if not, determine the rightful owner, and deny or clean-up the registration that is duplicate.

A.6. Requirements Related to Scalability

Use cases from Automatic Meter Reading (AMR, collection tree operations) and Advanced Metering Infrastructure (AMI, bi-directional communication to the meters) indicate the needs for a large number of LLN nodes pertaining to a single RPL DODAG (e.g. 5000) and connected to the 6LBR over a large number of LLN hops (e.g. 15).

Related requirements are:

Req6.1: The registration mechanism SHOULD enable a single 6LBR to register multiple thousands of devices.

Req6.2: The timing of the registration operation should allow for a large latency such as found in LLNs with ten and more hops.

Author's Address

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

