

6lo Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 26, 2017

C. Gomez
J. Paradells
UPC/i2CAT
J. Crowcroft
University of Cambridge
October 23, 2016

Optimized 6LoWPAN Fragmentation Header
draft-gomez-6lo-optimized-fragmentation-header-00

Abstract

RFC 4944 specifies 6LoWPAN fragmentation, in order to support the IPv6 MTU requirement over IEEE 802.15.4-2003 networks. The 6LoWPAN fragmentation header format comprises a 4-byte format for the first fragment, and a 5-byte format for subsequent fragments. This specification defines a more efficient 3-byte, optimized 6LoWPAN fragmentation header for all fragments.

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 April 26, 2017.

Copyright Notice

Copyright (c) 2016 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
1.1. Conventions used in this document	3
2. 6LoFH rules and format	3
3. Changes from RFC 4944 fragmentation header and rationale . .	4
4. IANA Considerations	5
5. Security Considerations	5
6. Acknowledgments	6
7. Annex A. Quantitative performance comparison of RFC 4944 fragmentation header with 6LoFH	7
8. References	7
8.1. Normative References	7
8.2. Informative References	8
Authors' Addresses	8

1. Introduction

IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) was originally designed as an adaptation layer intended to enable IPv6 over IEEE 802.15.4- 2003 networks [RFC4944]. One of the 6LoWPAN protocol suite components is fragmentation, which fulfills the IPv6 MTU requirement of 1280 bytes [RFC2460] over a radio interface with a layer two (L2) payload size around 100 bytes (in the best case) and without fragmentation support [RFC4944].

RFC 4944 defines the 6LoWPAN fragmentation header format, which comprises a 4-byte format for the first fragment, and a 5-byte format for subsequent fragments. This specification defines a more efficient 3-byte, optimized 6LoWPAN Fragmentation Header (6LoFH). The benefits of using 6LoFH are the following:

- Reduced overhead for transporting an IPv6 packet that requires fragmentation (see Annex A). This decreases consumption of energy and bandwidth, which are typically limited resources in the scenarios where 6LoWPAN fragmentation is used.

- Because the datagram offset can be expressed in increments of a single octet, 6LoFH enables the transport of IPv6 packets over L2 data units with a maximum payload size as small as only 4 bytes in the most extreme case. Note that RFC 4944 fragmentation can only be used over L2 technologies with a maximum L2 payload size of at least 13 bytes.

In comparison with the 6LoWPAN fragmentation header, parsing of the 6loFH format is also simplified, as the format has a constant size, and a 'symmetric' shape for both the first fragment and subsequent fragments. However, receiver buffer management will involve greater complexity as explained in Section 3.

1.1. Conventions used in this document

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]

2. 6LoFH rules and format

If an entire payload (e.g., IPv6) datagram fits within a single L2 data unit, it is unfragmented and a fragmentation header is not needed. If the datagram does not fit within a single L2 data unit, it SHALL be broken into fragments. The first fragment SHALL contain the first fragment header as defined in Figure 1.




Figure 1: First Fragment

The second and subsequent fragments (up to and including the last) SHALL contain a fragmentation header that conforms to the format shown in Figure 2.




Figure 2: Subsequent Fragments

datagram_size: This 11-bit field encodes the size of the entire IP packet before link-layer fragmentation (but after IP layer fragmentation). For IPv6, the datagram size SHALL be 40 octets (the size of the uncompressed IPv6 header) more than the value of Payload Length in the IPv6 header [RFC4944] of the packet. Note that this

packet may already be fragmented by hosts involved in the communication, i.e., this field needs to encode a maximum length of 1280 octets (the required by IPv6).

datagram_tag: The value of `datagram_tag` (datagram tag) SHALL be the same for all fragments of a payload (e.g., IPv6) datagram. The sender SHALL increment `datagram_tag` for successive, fragmented datagrams. The incremented value of `datagram_tag` SHALL wrap from 255 back to zero. This field is 8 bits long, and its initial value is not defined.

datagram_offset: This field is present only in the second and subsequent fragments and SHALL specify the offset, in increments of 1 octet, of the fragment from the beginning of the payload datagram. The first octet of the datagram (e.g., the start of the IPv6 header) has an offset of zero; the implicit value of `datagram_offset` in the first fragment is zero. This field is 11 bits long.

The recipient of link fragments SHALL use (1) the sender's L2 source address, (2) the destination's L2 address, (3) `datagram_size`, and (4) `datagram_tag` to identify all the fragments that belong to a given datagram.

Upon receipt of a link fragment, the recipient starts constructing the original unfragmented packet whose size is `datagram_size`. It uses the `datagram_offset` field to determine the location of the individual fragments within the original unfragmented packet. For example, it may place the data payload (except the encapsulation header) within a payload datagram reassembly buffer at the location specified by `datagram_offset`. The size of the reassembly buffer SHALL be determined from `datagram_size`.

If a fragment recipient disassociates from its L2 network, the recipient MUST discard all link fragments of all partially reassembled payload datagrams, and fragment senders MUST discard all not yet transmitted link fragments of all partially transmitted payload (e.g., IPv6) datagrams. Similarly, when a node first receives a fragment with a given `datagram_tag`, it starts a reassembly timer. When this time expires, if the entire packet has not been reassembled, the existing fragments MUST be discarded and the reassembly state MUST be flushed. The reassembly timeout MUST be set to a maximum of TBD seconds).

3. Changes from RFC 4944 fragmentation header and rationale

The main changes introduced in this specification to the fragmentation header format defined in RFC 4944 are listed below, together with their rationale:

-- The datagram size field is only included in the first fragment.
Rationale: In the RFC 4944 fragmentation header, the datagram size was included in all fragments to ease the task of reassembly at the receiver, since in an IEEE 802.15.4 mesh network, the fragment that arrives earliest to a destination is not necessarily the first fragment transmitted by the source. Nevertheless, the fragmentation format defined in this document supports reordering, at the expense of additional complexity in this regard.

-- The datagram tag size is reduced from 2 bytes to 1 byte.
Rationale: Given the low bit rate, as well as the relatively low message rate in IEEE 802.15.4 scenarios, ambiguities due to datagram tag wrapping events are unlikely despite the reduced tag space.

-- The datagram offset size is increased from 8 bits to 11 bits.
Rationale: This allows to express the datagram offset in single-octet increments.

4. IANA Considerations

This document allocates the following sixteen RFC 4944 Dispatch type values:

11001 000

through

11001 111

and

11010 000

through

11010 111

5. Security Considerations

6LoWPAN fragmentation attacks have been analyzed in the literature. Countermeasures to these have been proposed as well [HHWH].

A node can perform a buffer reservation attack by sending a first fragment to a target. Then, the receiver will reserve buffer space for the whole packet on the basis of the datagram size announced in that first fragment. Other incoming fragmented packets will be dropped while the reassembly buffer is occupied during the reassembly timeout. Once that timeout expires, the attacker can repeat the same

procedure, and iterate, thus creating a denial of service attack. The (low) cost to mount this attack is linear with the number of buffers at the target node. However, the cost for an attacker can be increased if individual fragments of multiple packets can be stored in the reassembly buffer. To further increase the attack cost, the reassembly buffer can be split into fragment-sized buffer slots. Once a packet is complete, it is processed normally. If buffer overload occurs, a receiver can discard packets based on the sender behavior, which may help identify which fragments have been sent by an attacker.

In another type of attack, the malicious node is required to have overhearing capabilities. If an attacker can overhear a fragment, it can send a spoofed duplicate (e.g. with random payload) to the destination. A receiver cannot distinguish legitimate from spoofed fragments. Therefore, the original IPv6 packet will be considered corrupt and will be dropped. To protect resource-constrained nodes from this attack, it has been proposed to establish a binding among the fragments to be transmitted by a node, by applying content-chaining to the different fragments, based on cryptographic hash functionality. The aim of this technique is to allow a receiver to identify illegitimate fragments.

Further attacks may involve sending overlapped fragments (i.e. comprising some overlapping parts of the original datagram) or announcing a datagram size in the first fragment that does not reflect the actual amount of data carried by the fragments. Implementers should make sure that correct operation is not affected by such events.

6. Acknowledgments

In section 2, the authors have reused extensive parts of text available in section 5.3 of RFC 4944, and would like to thank the authors of RFC 4944.

The authors would like to thank Carsten Bormann, Tom Phinney, Ana Minaburo and Laurent Toutain for valuable comments that helped improve the document.

Carles Gomez has been funded in part by the Spanish Government (Ministerio de Educacion, Cultura y Deporte) through the Jose Castillejo grant CAS15/00336. Part of his contribution to this work has been carried out during his stay as a visiting scholar at the Computer Laboratory of the University of Cambridge.

7. Annex A. Quantitative performance comparison of RFC 4944 fragmentation header with 6LoFH

	IPv6 datagram size (bytes)							
	40		100		640		1280	
L2 payload (bytes)	4944	6LoFH	4944	6LoFH	4944	6LoFH	4944	6LoFH
10	----	18	----	45	----	276	----	549
20	19	9	59	18	394	114	794	228
40	0	0	19	9	99	54	199	105
60	0	0	9	6	69	36	134	69
80	0	0	9	6	44	27	89	51
100	0	0	0	0	39	21	74	42

Figure 3: Adaptation layer fragmentation overhead (in bytes) required to transport an IPv6 datagram

Note 1: while IEEE 802.15.4-2003 allows a maximum L2 payload size between 81 and 102 bytes, a range of L2 payload size between 10 and 100 bytes is considered in the study to illustrate the performance of 6LoFH also for other potential L2 technologies with short payload size and without fragmentation support.

Note 2: with the RFC 4944 fragmentation header it is not possible to transport IPv6 datagrams of the considered sizes over a 10-byte payload L2 technology.

8. References

8.1. 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>>.

- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<http://www.rfc-editor.org/info/rfc4944>>.
- [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>>.

8.2. Informative References

- [HHWH] Hummen et al, R., "6LoWPAN fragmentation attacks and mitigation mechanisms", 2013.
- [I-D.minaburo-lpwan-gap-analysis] Minaburo, A., Gomez, C., Toutain, L., Paradells, J., and J. Crowcroft, "LPWAN Survey and GAP Analysis", draft-minaburo-lpwan-gap-analysis-02 (work in progress), October 2016.

Authors' Addresses

Carles Gomez
UPC/i2CAT
C/Esteve Terradas, 7
Castelldefels 08860
Spain

Email: carlesgo@entel.upc.edu

Josep Paradells
UPC/i2CAT
C/Jordi Girona, 1-3
Barcelona 08034
Spain

Email: josep.paradells@entel.upc.edu

Jon Crowcroft
University of Cambridge
JJ Thomson Avenue
Cambridge, CB3 0FD
United Kingdom

Email: jon.crowcroft@cl.cam.ac.uk

6lo
Internet-Draft
Updates: 8505 (if approved)
Intended status: Standards Track
Expires: 1 November 2020

P. Thubert, Ed.
Cisco
B. Sarikaya

M. Sethi
Ericsson
R. Struik
Struik Security Consultancy
30 April 2020

Address Protected Neighbor Discovery for Low-power and Lossy Networks
draft-ietf-6lo-ap-nd-23

Abstract

This document updates the 6LoWPAN Neighbor Discovery (ND) protocol defined in RFC 6775 and RFC 8505. The new extension is called Address Protected Neighbor Discovery (AP-ND) and it protects the owner of an address against address theft and impersonation attacks in a low-power and lossy network (LLN). Nodes supporting this extension compute a cryptographic identifier (Crypto-ID) and use it with one or more of their Registered Addresses. The Crypto-ID identifies the owner of the Registered Address and can be used to provide proof of ownership of the Registered Addresses. Once an address is registered with the Crypto-ID and a proof-of-ownership is provided, only the owner of that address can modify the registration information, thereby enforcing Source Address Validation.

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 1 November 2020.

Copyright Notice

Copyright (c) 2020 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	3
2. Terminology	4
2.1. BCP 14	4
2.2. Additional References	4
2.3. Abbreviations	5
3. Updating RFC 8505	5
4. New Fields and Options	6
4.1. New Crypto-ID	6
4.2. Updated EARO	7
4.3. Crypto-ID Parameters Option	8
4.4. NDP Signature Option	10
4.5. Extensions to the Capability Indication Option	11
5. Protocol Scope	12
6. Protocol Flows	13
6.1. First Exchange with a 6LR	14
6.2. NDPSO generation and verification	16
6.3. Multihop Operation	17
7. Security Considerations	18
7.1. Brown Field	18
7.2. Inheriting from RFC 3971	18
7.3. Related to 6LoWPAN ND	19
7.4. Compromised 6LR	20
7.5. ROVR Collisions	20
7.6. Implementation Attacks	21
7.7. Cross-Algorithm and Cross-Protocol Attacks	21
7.8. Public Key Validation	22
7.9. Correlating Registrations	22
8. IANA considerations	22
8.1. CGA Message Type	22
8.2. Crypto-Type Subregistry	23
8.3. IPv6 ND option types	24
8.4. New 6LoWPAN Capability Bit	24

9. Acknowledgments	24
10. Normative References	24
11. Informative references	26
Appendix A. Requirements Addressed in this Document	28
Appendix B. Representation Conventions	28
B.1. Signature Schemes	28
B.2. Representation of ECDSA Signatures	29
B.3. Representation of Public Keys Used with ECDSA	30
B.4. Alternative Representations of Curve25519	30
Authors' Addresses	32

1. Introduction

Neighbor Discovery Optimizations for 6LoWPAN networks [RFC6775] (6LoWPAN ND) adapts the original IPv6 Neighbor Discovery (IPv6 ND) protocols defined in [RFC4861] and [RFC4862] for constrained low-power and lossy network (LLN). In particular, 6LoWPAN ND introduces a unicast host Address Registration mechanism that reduces the use of multicast compared to the Duplicate Address Detection (DAD) mechanism defined in IPv6 ND. 6LoWPAN ND defines a new Address Registration Option (ARO) that is carried in the unicast Neighbor Solicitation (NS) and Neighbor Advertisement (NA) messages exchanged between a 6LoWPAN Node (6LN) and a 6LoWPAN Router (6LR). It also defines the Duplicate Address Request (DAR) and Duplicate Address Confirmation (DAC) messages between the 6LR and the 6LoWPAN Border Router (6LBR). In LLN networks, the 6LBR is the central repository of all the registered addresses in its domain.

The registration mechanism in "Neighbor Discovery Optimization for Low-power and Lossy Networks" [RFC6775] (aka 6LoWPAN ND) prevents the use of an address if that address is already registered in the subnet (first come first serve). In order to validate address ownership, the registration mechanism enables the 6LR and 6LBR to validate the association between the registered address of a node, and its Registration Ownership Verifier (ROVR). The ROVR is defined in "Registration Extensions for 6LoWPAN Neighbor Discovery" [RFC8505] and it can be derived from the MAC address of the device (using the 64-bit Extended Unique Identifier EUI-64 address format specified by IEEE). However, the EUI-64 can be spoofed, and therefore, any node connected to the subnet and aware of a registered-address-to-ROVR mapping could effectively fake the ROVR. This would allow an attacker to steal the address and redirect traffic for that address. [RFC8505] defines an Extended Address Registration Option (EARO) option that transports alternate forms of ROVRs, and is a pre-requisite for this specification.

In this specification, a 6LN generates a cryptographic ID (Crypto-ID) and places it in the ROVR field during the registration of one (or more) of its addresses with the 6LR(s). Proof of ownership of the Crypto-ID is passed with the first registration exchange to a new 6LR, and enforced at the 6LR. The 6LR validates ownership of the cryptographic ID before it creates any new registration state, or changes existing information.

The protected address registration protocol proposed in this document provides the same conceptual benefit as Source Address Validation (SAVI) [RFC7039] that only the owner of an IPv6 address may source packets with that address. As opposed to [RFC7039], which relies on snooping protocols, the protection is based on a state that is installed and maintained in the network by the owner of the address. With this specification, a 6LN may use a 6LR for forwarding an IPv6 packets if and only if it has registered the address used as source of the packet with that 6LR.

With the 6lo adaptation layer in [RFC4944] and [RFC6282], a 6LN can obtain a better compression for an IPv6 address with an Interface ID (IID) that is derived from a Layer-2 address. As a side note, this is incompatible with Secure Neighbor Discovery (SeND) [RFC3971] and Cryptographically Generated Addresses (CGAs) [RFC3972], since they derive the IID from cryptographic keys, whereas this specification separates the IID and the key material.

2. Terminology

2.1. BCP 14

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2.2. Additional References

The reader may get additional context for this specification from the following references:

- * "SEcure Neighbor Discovery (SEND)" [RFC3971],
- * "Cryptographically Generated Addresses (CGA)" [RFC3972],
- * "Neighbor Discovery for IP version 6" [RFC4861] ,
- * "IPv6 Stateless Address Autoconfiguration" [RFC4862], and
- * "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals " [RFC4919].

2.3. Abbreviations

This document uses the following abbreviations:

6BBR: 6LoWPAN Backbone Router
6LBR: 6LoWPAN Border Router
6LN: 6LoWPAN Node
6LR: 6LoWPAN Router
CGA: Cryptographically Generated Address
EARO: Extended Address Registration Option
ECDH: Elliptic curve Diffie-Hellman
ECDSA: Elliptic Curve Digital Signature Algorithm
CIPO: Crypto-ID Parameters Option
LLN: Low-Power and Lossy Network
JSON: JavaScript Object Notation
JOSE: JavaScript Object Signing and Encryption
JWK: JSON Web Key
JWS: JSON Web Signature
NA: Neighbor Advertisement
ND: Neighbor Discovery
NDP: Neighbor Discovery Protocol
NDPSO: Neighbor Discovery Protocol Signature Option
NS: Neighbor Solicitation
ROVR: Registration Ownership Verifier
RA: Router Advertisement
RS: Router Solicitation
RSAO: RSA Signature Option
SHA: Secure Hash Algorithm
SLAAC: Stateless Address Autoconfiguration
TID: Transaction ID

3. Updating RFC 8505

Section 5.3 of [RFC8505] introduces the ROVR that is used to detect and reject duplicate registrations in the DAD process. The ROVR is a generic object that is designed for both backward compatibility and the capability to introduce new computation methods in the future. Using a Crypto-ID per this specification is the RECOMMENDED method. Section 7.5 discusses collisions when heterogeneous methods to compute the ROVR field coexist inside a same network.

This specification introduces a new token called a cryptographic identifier (Crypto-ID) that is transported in the ROVR field and used to prove indirectly the ownership of an address that is being registered by means of [RFC8505]. The Crypto-ID is derived from a cryptographic public key and additional parameters.

The overall mechanism requires the support of Elliptic Curve Cryptography (ECC) and of a hash function as detailed in Section 6.2. To enable the verification of the proof, the registering node needs to supply certain parameters including a nonce and a signature that will demonstrate that the node possesses the private-key corresponding to the public-key used to build the Crypto-ID.

The elliptic curves and the hash functions listed in Table 1 in Section 8.2 can be used with this specification; more may be added in the future to the IANA registry. The signature scheme that specifies which combination is used (including the curve and the representation conventions) is signaled by a Crypto-Type in a new IPv6 ND Crypto-ID Parameters Option (CIPO, see Section 4.3) that contains the parameters that are necessary for the proof, a Nonce option ([RFC3971]) and a NDP Signature option (Section 4.4). The NA(EARO) is modified to enable a challenge and transport a Nonce option.

4. New Fields and Options

4.1. New Crypto-ID

The Crypto-ID is transported in the ROVR field of the EARO option and the EDAR message, and is associated with the Registered Address at the 6LR and the 6LBR. The ownership of a Crypto-ID can be demonstrated by cryptographic mechanisms, and by association, the ownership of the Registered Address can be ascertained.

A node in possession of the necessary cryptographic primitives SHOULD use Crypto-ID by default as ROVR in its registrations. Whether a ROVR is a Crypto-ID is indicated by a new "C" flag in the NS(EARO) message.

The Crypto-ID is derived from the public key and a modifier as follows:

1. The hash function used internally by the signature scheme indicated by the Crypto-Type (see also Table 1 in Section 8.2) is applied to the CIPO. Note that all the reserved and padding bits MUST be set to zero.
2. The leftmost bits of the resulting hash, up to the desired size, are used as the Crypto-ID.

At the time of this writing, a minimal size for the Crypto-ID of 128 bits is RECOMMENDED unless backward compatibility is needed [RFC8505]. This value is bound to augment in the future.

4.2. Updated EARO

This specification updates the EARO option to enable the use of the ROVR field to transport the Crypto-ID. The resulting format is as follows:

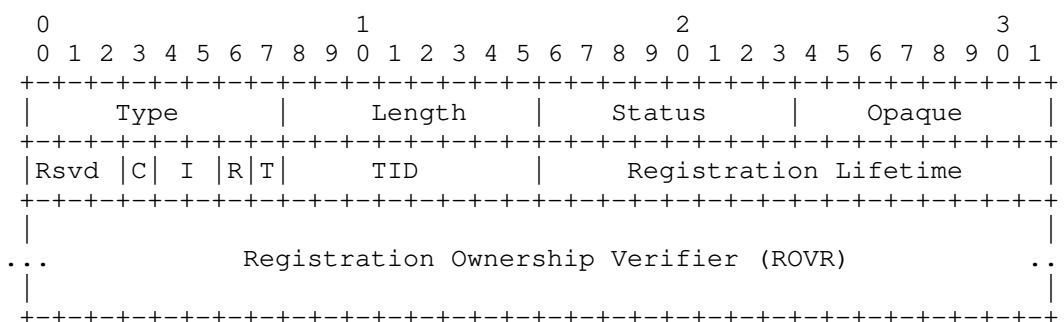


Figure 1: Enhanced Address Registration Option

Type: 33

Length: Defined in [RFC8505] and copied in associated CIPO.

Status: Defined in [RFC8505].

Opaque: Defined in [RFC8505].

Rsrd (Reserved): 3-bit unsigned integer. It MUST be set to zero by the sender and MUST be ignored by the receiver.

C: This "C" flag is set to indicate that the ROVR field contains a Crypto-ID and that the 6LN MAY be challenged for ownership as specified in this document.

I, R, T: Defined in [RFC8505].

TID: Defined in [RFC8505].

Registration Ownership Verifier (ROVR): When the "C" flag is set, this field contains a Crypto-ID.

This specification uses Status values "Validation Requested" and "Validation Failed", which are defined in [RFC8505].

this specification does not define any new Status value.

4.3. Crypto-ID Parameters Option

This specification defines the Crypto-ID Parameters Option (CIPO). The CIPO carries the parameters used to form a Crypto-ID.

In order to provide cryptographic agility [BCP 201], this specification supports different elliptic-curve based signature schemes, indicated by a Crypto-Type field:

- * The ECDSA256 signature scheme, which uses ECDSA with the NIST P-256 curve [FIPS186-4] and the hash function SHA-256 [RFC6234] internally, **MUST** be supported by all implementations.
- * The Ed25519 signature scheme, which uses the Pure Edwards-Curve Digital Signature Algorithm (PureEdDSA) [RFC8032] with the twisted Edwards curve Edwards25519 [RFC7748] and the hash function SHA-512 [RFC6234] internally, **MAY** be supported as an alternative.
- * The ECDSA25519 signature scheme, which uses ECDSA [FIPS186-4] with the Weierstrass curve Wei25519 (see Appendix B.4) and the hash function SHA-256 [RFC6234] internally, **MAY** also be supported.

This specification uses signature schemes that target similar cryptographic strength but rely on different curves, hash functions, signature algorithms, and/or representation conventions. Future specification may extend this to different cryptographic algorithms and key sizes, e.g., to provide better security properties or a simpler implementation.




Figure 2: Crypto-ID Parameters Option

Type: 8-bit unsigned integer. to be assigned by IANA, see Table 2.

Length: 8-bit unsigned integer. The length of the option in units of 8 octets.

Reserved1: 5-bit unsigned integer. It MUST be set to zero by the sender and MUST be ignored by the receiver.

Public Key Length: 11-bit unsigned integer. The length of the Public Key field in bytes. The actual length depends on the Crypto-Type value and on how the public key is represented. The valid values with this document are provided in Table 1.

Crypto-Type: 8-bit unsigned integer. The type of cryptographic algorithm used in calculation Crypto-ID indexed by IANA in the "Crypto-Type Subregistry" in the "Internet Control Message Protocol version 6 (ICMPv6) Parameters" (see Section 8.2).

Modifier: 8-bit unsigned integer. Set to an arbitrary value by the creator of the Crypto-ID. The role of the modifier is to enable the formation of multiple Crypto-IDs from a same key pair, which reduces the traceability and thus improves the privacy of a constrained node that could not maintain many key-pairs.

EARO Length: 8-bit unsigned integer. The option length of the EARO that contains the Crypto-ID associated with the CIPO.

Public Key: A variable-length field, size indicated in the Public Key Length field.

Padding: A variable-length field completing the Public Key field to align to the next 8-bytes boundary. It MUST be set to zero by the sender and MUST be ignored by the receiver.

The implementation of multiple hash functions in a constrained device may consume excessive amounts of program memory. This specification enables the use of the same hash function SHA-256 [RFC6234] for two of the three supported ECC-based signature schemes. Some code factorization is also possible for the ECC computation itself.

[CURVE-REPR] provides information on how to represent Montgomery curves and (twisted) Edwards curves as curves in short-Weierstrass form and illustrates how this can be used to implement elliptic curve computations using existing implementations that already provide, e.g., ECDSA and ECDH using NIST [FIPS186-4] prime curves. For more details on representation conventions, we refer to Appendix B.

4.4. NDP Signature Option

This specification defines the NDP Signature Option (NDPSO). The NDPSO carries the signature that proves the ownership of the Crypto-ID. The format of the NDPSO is illustrated in Figure 3.

As opposed to the RSA Signature Option (RSAO) defined in section 5.2. of SEND [RFC3971], the NDPSO does not have a key hash field. Instead, the leftmost 128 bits of the ROVR field in the EARO are used as hash to retrieve the CIPO that contains the key material used for signature verification, left-padded if needed.

Another difference is that the NDPSO signs a fixed set of fields as opposed to all options that appear prior to it in the ND message that bears the signature. This allows to elide a CIPO that the 6LR already received, at the expense of the capability to add arbitrary options that would signed with a RSAO.

An ND message that carries an NDPSO MUST have one and only one EARO. The EARO MUST contain a Crypto-ID in the ROVR field, and the Crypto-ID MUST be associated with the keypair used for the Digital Signature in the NDPSO.

The CIPO may be present in the same message as the NDPSO. If it is not present, it can be found in an abstract table that was created by a previous message and indexed by the hash.




Figure 3: NDP signature Option

Type: to be assigned by IANA, see Table 2.

Length: 8-bit unsigned integer. The length of the option in units of 8 octets.

Reserved1: 5-bit unsigned integer. It MUST be set to zero by the sender and MUST be ignored by the receiver.

Digital Signature Length: 11-bit unsigned integer. The length of the Digital Signature field in bytes.

Reserved2: 32-bit unsigned integer. It MUST be set to zero by the sender and MUST be ignored by the receiver.

Digital Signature: A variable-length field containing the digital signature. The length and computation of the digital signature both depend on the Crypto-Type which is found in the associated CIPO, see Appendix B. For the values of the Crypto-Type defined in this specification, and for future values of the Crypto-Type unless specified otherwise, the signature is computed as detailed in Section 6.2.

Padding: A variable-length field completing the Digital Signature field to align to the next 8-bytes boundary. It MUST be set to zero by the sender and MUST be ignored by the receiver.

4.5. Extensions to the Capability Indication Option

This specification defines one new capability bits in the 6CIO, defined by [RFC7400] for use by the 6LR and 6LBR in IPv6 ND RA messages.




Figure 4: New Capability Bit in the 6CIO

New Option Field:

A: 1-bit flag. Set to indicate that AP-ND is globally activated in the network.

6. Protocol Flows

The 6LR/6LBR ensures first-come/first-serve by storing the ROVR associated to the address being registered upon the first registration and rejecting a registration with a different ROVR value. A 6LN can claim any address as long as it is the first to make that claim. After a successful registration, the 6LN becomes the owner of the registered address and the address is bound to the ROVR value in the 6LR/6LBR registry.

This specification protects the ownership of the address at the first hop (the edge). Its use in a network is signaled by the "A" flag in the 6CIO. The flag is set by the 6LBR and propagated unchanged by the 6LRs. The "A" flag enables to migrate a network with the protection off and then turn it on globally.

The 6LN places a cryptographic token, the Crypto-ID, in the ROVR that is associated with the address at the first registration, enabling the 6LR to later challenge it to verify that it is the original Registering Node. The challenge may happen at any time at the discretion of the 6LR and the 6LBR. A valid registration in the 6LR or the 6LBR MUST NOT be altered until the challenge is complete.

When the "A" flag is on, the 6LR MUST challenge the 6LN when it creates a binding with the "C" flag set in the ROVR and when a new registration attempts to change a parameter of that binding that identifies the 6LN, for instance its Source Link-Layer Address. The verification protects against a rogue that would steal an address and attract its traffic, or use it as source address.

The 6LR MUST also challenge the 6LN if the 6LBR directly signals to do so, using an EDAC Message with a "Validation Requested" status. The EDAR is echoed by the 6LR in the NA (EARO) back to the registering node. The 6LR SHOULD also challenge all its attached 6LNs at the time the 6LBR turns the "A" flag on in the 6CIO, to detect an issue immediately.

If the 6LR does not support the Crypto-Type, it MUST reply with an EARO Status 10 "Validation Failed" without a challenge. In that case, the 6LN may try another Crypto-Type until it falls back to Crypto-Type 0 that MUST be supported by all 6LRs.

A node may use more than one IPv6 address at the same time. The separation of the address and the cryptographic material avoids the need for the constrained device to compute multiple keys for multiple addresses. The 6LN MAY use the same Crypto-ID to prove the ownership of multiple IPv6 addresses. The 6LN MAY also derive multiple Crypto-IDs from a same key.

6.1. First Exchange with a 6LR

A 6LN registers to a 6LR that is one hop away from it with the "C" flag set in the EARO, indicating that the ROVR field contains a Crypto-ID. The Target Address in the NS message indicates the IPv6 address that the 6LN is trying to register [RFC8505]. The on-link (local) protocol interactions are shown in Figure 6. If the 6LR does not have a state with the 6LN that is consistent with the NS (EARO), then it replies with a challenge NA (EARO, status=Validation Requested) that contains a Nonce Option (shown as NonceLR in Figure 6).

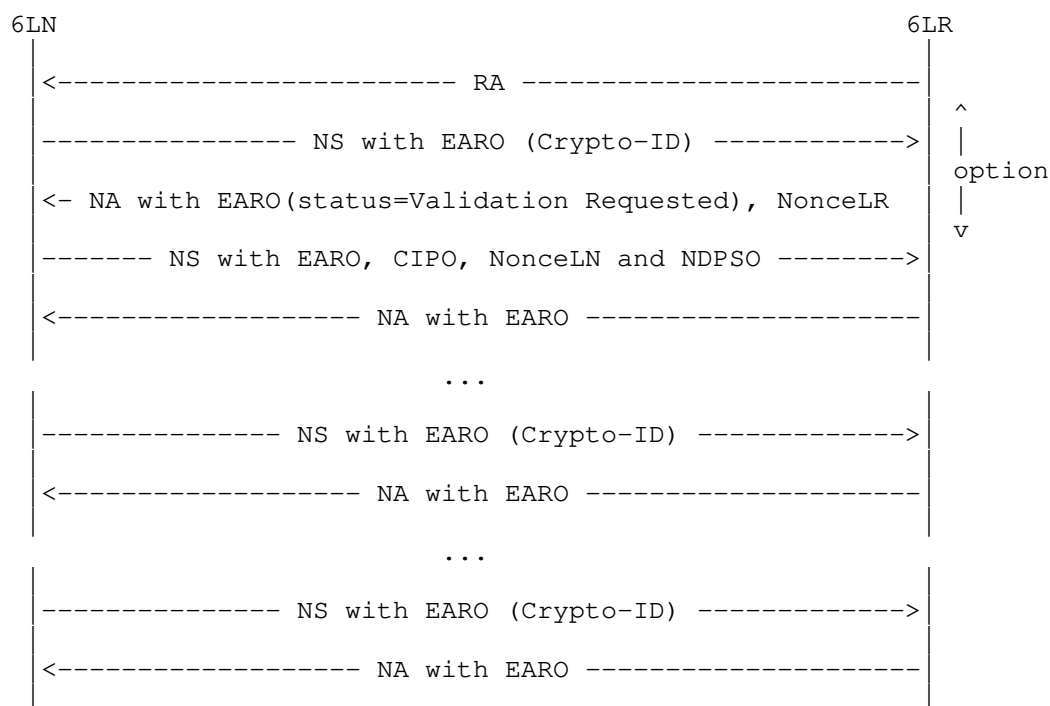


Figure 6: On-link Protocol Operation

The Nonce option contains a nonce value that, to the extent possible for the implementation, was never employed in association with the key pair used to generate the Crypto-ID. This specification inherits from [RFC3971] that simply indicates that the nonce is a random value. Ideally, an implementation uses an unpredictable cryptographically random value [BCP 106]. But that may be impractical in some LLN scenarios where the devices do not have a guaranteed sense of time and for which computing complex hashes is detrimental to the battery lifetime.

Alternatively, the device may use an always-incrementing value saved in the same stable storage as the key, so they are lost together, and starting at a best effort random value, either as the nonce value or as a component to its computation.

The 6LN replies to the challenge with an NS(EARO) that includes a new Nonce option (shown as NonceLN in Figure 6), the CIPO (Section 4.3), and the NDPSO containing the signature. Both Nonces are included in the signed material. This provides a "contributory behavior", so that either party that knows it generates a good quality nonce knows that the protocol will be secure.

The 6LR MUST store the information associated to a Crypto-ID on the first NS exchange where it appears in a fashion that the CIPO parameters can be retrieved from the Crypto-ID alone.

The steps for the registration to the 6LR are as follows:

Upon the first exchange with a 6LR, a 6LN will be challenged to prove ownership of the Crypto-ID and the Target Address being registered in the Neighbor Solicitation message. When a 6LR receives a NS(EARO) registration with a new Crypto-ID as a ROVR, and unless the registration is rejected for another reason, it MUST challenge by responding with a NA(EARO) with a status of "Validation Requested".

Upon receiving a first NA(EARO) with a status of "Validation Requested" from a 6LR, the registering node SHOULD retry its registration with a Crypto-ID Parameters Option (CIPO) (Section 4.3) that contains all the necessary material for building the Crypto-ID, the NonceLN that it generated, and the NDP signature (Section 4.4) option that proves its ownership of the Crypto-ID and intent of registering the Target Address. In subsequent revalidation with the same 6LR, the 6LN MAY try to omit the CIPO to save bandwidth, with the expectation that the 6LR saved it. If the validation fails and it gets challenged again, then it SHOULD add the CIPO again.

In order to validate the ownership, the 6LR performs the same steps as the 6LN and rebuilds the Crypto-ID based on the parameters in the CIPO. If the rebuilt Crypto-ID matches the ROVR, the 6LR also verifies the signature contained in the NDPSO option. If at that point the signature in the NDPSO option can be verified, then the validation succeeds. Otherwise the validation fails.

If the 6LR fails to validate the signed NS(EARO), it responds with a status of "Validation Failed". After receiving a NA(EARO) with a status of "Validation Failed", the registering node SHOULD try and alternate Crypto-Type and if even Crypto-Type 0 fails, it may try to register a different address in the NS message.

6.2. NDPSO generation and verification

The signature generated by the 6LN to provide proof-of-ownership of the private-key is carried in the NDP Signature Option (NDPSO). It is generated by the 6LN in a fashion that depends on the Crypto-Type (see Table 1 in Section 8.2) chosen by the 6LN as follows:

- * Form the message to be signed, by concatenating the following byte-strings in the order listed:
 1. The 128-bit Message Type tag [RFC3972] (in network byte order). For this specification the tag is given in Section 8.1. (The tag value has been generated by the editor of this specification on random.org).
 2. the CIPO
 3. the 16-byte Target Address (in network byte order) sent in the Neighbor Solicitation (NS) message. It is the address which the 6LN is registering with the 6LR and 6LBR.
 4. NonceLR received from the 6LR (in network byte order) in the Neighbor Advertisement (NA) message. The nonce is at least 6 bytes long as defined in [RFC3971].
 5. NonceLN sent from the 6LN (in network byte order). The nonce is at least 6 bytes long as defined in [RFC3971].
 6. 1-byte Option Length of the EARO containing the Crypto-ID.
- * Apply the signature algorithm specified by the Crypto-Type using the private key.

The 6LR on receiving the NDPSO and CIPO options first checks that the EARO Length in the CIPO matches the length of the EARO. If so it regenerates the Crypto-ID based on the CIPO to make sure that the leftmost bits up to the size of the ROVR match.

If and only if the check is successful, it tries to verify the signature in the NDPSO option using the following:

- * Form the message to be verified, by concatenating the following byte-strings in the order listed:
 1. The 128-bit Message Type tag given in Section 8.1 (in network byte order)
 2. the CIPO
 3. the 16-byte Target Address (in network byte order) received in the Neighbor Solicitation (NS) message. It is the address which the 6LN is registering with the 6LR and 6LBR.
 4. NonceLR sent in the Neighbor Advertisement (NA) message. The nonce is at least 6 bytes long as defined in [RFC3971].

5. NonceLN received from the 6LN (in network byte order) in the NS message. The nonce is at least 6 bytes long as defined in [RFC3971].
 6. 1-byte EARO Length received in the CIP0.
- * Verify the signature on this message with the public-key in the CIP0 and the locally computed values using the signature algorithm specified by the Crypto-Type. If the verification succeeds, the 6LR propagates the information to the 6LBR using a EDAR/EDAC flow.
 - * Due to the first-come/first-serve nature of the registration, if the address is not registered to the 6LBR, then flow succeeds and both the 6LR and 6LBR add the state information about the Crypto-ID and Target Address being registered to their respective abstract database.

6.3. Multihop Operation

A new 6LN that joins the network auto-configures an address and performs an initial registration to a neighboring 6LR with an NS message that carries an Address Registration Option (EARO) [RFC8505].

In a multihop 6LoWPAN, the registration with Crypto-ID is propagated to 6LBR as shown in Figure 7, which illustrates the registration flow all the way to a 6LoWPAN Backbone Router (6BBR) [BACKBONE-ROUTER].

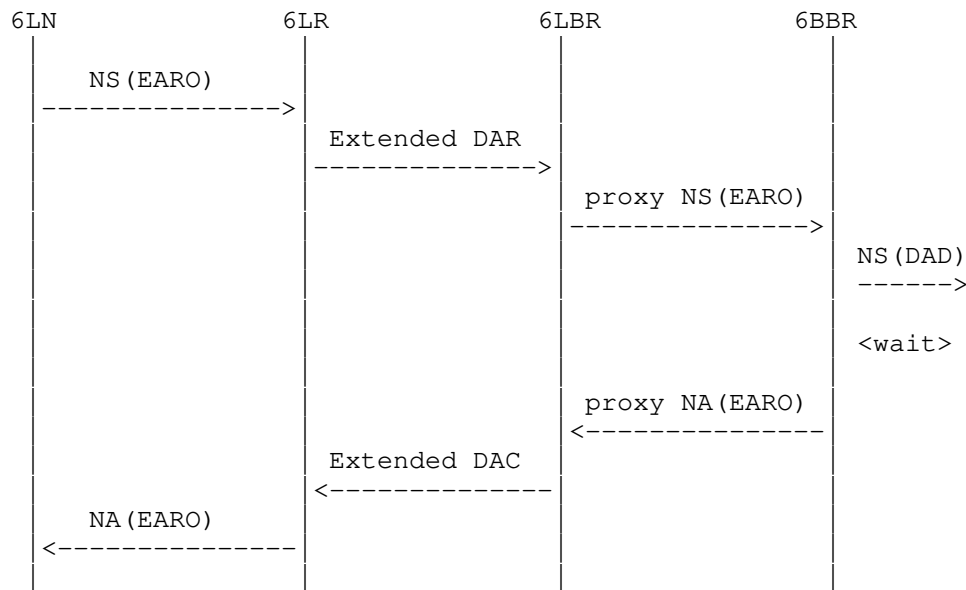


Figure 7: (Re-)Registration Flow

The 6LR and the 6LBR communicate using ICMPv6 Extended Duplicate Address Request (EDAR) and Extended Duplicate Address Confirmation (EDAC) messages [RFC8505] as shown in Figure 7. This specification extends EDAR/EDAC messages to carry cryptographically generated ROVR.

The assumption is that the 6LR and the 6LBR maintain a security association to authenticate and protect the integrity of the EDAR and EDAC messages, so there is no need to propagate the proof of ownership to the 6LBR. The 6LBR implicitly trusts that the 6LR performs the verification when the 6LBR requires it, and if there is no further exchange from the 6LR to remove the state, that the verification succeeded.

7. Security Considerations

7.1. Brown Field

Only 6LRs that are upgraded to this specification are capable to challenge a registration and repel an attack. In a brown (mixed) network, an attacker may attach to a legacy 6LR and fool the 6LBR. So even if the "A" flag could be set at any time to test the protocol operation, the security will only be effective when all the 6LRs are upgraded.

7.2. Inheriting from RFC 3971

Observations regarding the following threats to the local network in [RFC3971] also apply to this specification.

Neighbor Solicitation/Advertisement Spoofing: Threats in section 9.2.1 of RFC3971 apply. AP-ND counters the threats on NS(EARO) messages by requiring that the NDP Signature and CIPO options be present in these solicitations.

Duplicate Address Detection DoS Attack: Inside the LLN, Duplicate Addresses are sorted out using the ROVR, which differentiates it from a movement. A different ROVR for the same Registered address entails a rejection of the second registration [RFC8505]. DAD coming from the backbone are not forwarded over the LLN, which provides some protection against DoS attacks inside the resource-constrained part of the network. Over the backbone, the EARO option is present in NS/NA messages. This protects against misinterpreting a movement for a duplication, and enables the backbone routers to determine which one has the freshest registration [RFC8505] and is thus the best candidate to validate the registration for the device attached to it [BACKBONE-ROUTER]. But this specification does not guarantee that the backbone router claiming an address over the backbone is not an attacker.

Router Solicitation and Advertisement Attacks: This specification does not change the protection of RS and RA which can still be protected by SEND.

Replay Attacks A nonce should never repeat for a single key, but nonces do not need to be unpredictable for secure operation. Using nonces (NonceLR and NonceLN) generated by both the 6LR and 6LN ensure a contributory behavior that provides an efficient protection against replay attacks of the challenge/response flow. The quality of the protection by a random nonce depends on the random number generator and its parameters (e.g., sense of time).

Neighbor Discovery DoS Attack: A rogue node that managed to access the L2 network may form many addresses and register them using AP-ND. The perimeter of the attack is all the 6LRs in range of the attacker. The 6LR MUST protect itself against overflows and reject excessive registration with a status 2 "Neighbor Cache Full". This effectively blocks another (honest) 6LN from registering to the same 6LR, but the 6LN may register to other 6LRs that are in its range but not in that of the rogue.

7.3. Related to 6LoWPAN ND

The threats and mediations discussed in 6LoWPAN ND [RFC6775][RFC8505] also apply here, in particular denial-of-service attacks against the registry at the 6LR or 6LBR.

Secure ND [RFC3971] forces the IPv6 address to be cryptographic since it integrates the CGA as the IID in the IPv6 address. In contrast, this specification saves about 1Kbyte in every NS/NA message. Also, this specification separates the cryptographic identifier from the registered IPv6 address so that a node can have more than one IPv6 address protected by the same cryptographic identifier.

With this specification the 6LN can freely form its IPv6 address(es) in any fashion, thereby enabling either 6LoWPAN compression for IPv6 addresses that are derived from Layer-2 addresses, or temporary addresses, e.g., formed pseudo-randomly and released in relatively short cycles for privacy reasons [RFC8064][RFC8065], that cannot be compressed.

This specification provides added protection for addresses that are obtained following due procedure [RFC8505] but does not constrain the way the addresses are formed or the number of addresses that are used in parallel by a same entity. A rogue may still perform denial-of-service attack against the registry at the 6LR or 6LBR, or attempt to deplete the pool of available addresses at Layer-2 or Layer-3.

7.4. Compromised 6LR

This specification distributes the challenge and its validation at the edge of the network, between the 6LN and its 6LR. This protects against DOS attacks targeted at that central 6LBR. This also saves back and forth exchanges across a potentially large and constrained network.

The downside is that the 6LBR needs to trust the 6LR for performing the checking adequately, and the communication between the 6LR and the 6LBR must be protected to avoid tampering with the result of the test.

If a 6LR is compromised, and provided that it knows the ROVR field used by the real owner of the address, the 6LR may pretend that the owner has moved, is now attached to it and has successfully passed the Crpto-ID validation. The 6LR may then attract and inject traffic at will on behalf of that address or let a rogue take ownership of the address.

7.5. ROVR Collisions

A collision of Registration Ownership Verifiers (ROVR) (i.e., the Crypto-ID in this specification) is possible, but it is a rare event. Assuming in the calculations/discussion below that the hash used for calculating the Crypto-ID is a well-behaved cryptographic hash and thus that random collisions are the only ones possible, the formula (birthday paradox) for calculating the probability of a collision is $1 - e^{-p^2/(2n)}$ where n is the maximum population size (2^{64} here, 1.84E19) and p is the actual population (number of nodes, assuming one Crypto-ID per node).

If the Crypto-ID is 64-bits (the least possible size allowed), the chance of a collision is 0.01% for network of 66 million nodes. Moreover, the collision is only relevant when this happens within one stub network (6LBR). In the case of such a collision, a third party node would be able to claim the registered address of an another legitimate node, provided that it wishes to use the same address. To prevent address disclosure and avoid the chances of collision on both the ROVR and the address, it is RECOMMENDED that nodes do not derive the address being registered from the ROVR.

7.6. Implementation Attacks

The signature schemes referenced in this specification comply with NIST [FIPS186-4] or Crypto Forum Research Group (CFRG) standards [RFC8032] and offer strong algorithmic security at roughly 128-bit security level. These signature schemes use elliptic curves that were either specifically designed with exception-free and constant-time arithmetic in mind [RFC7748] or where one has extensive implementation experience of resistance to timing attacks [FIPS186-4].

However, careless implementations of the signing operations could nevertheless leak information on private keys. For example, there are micro-architectural side channel attacks that implementors should be aware of [breaking-ed25519]. Implementors should be particularly aware that a secure implementation of Ed25519 requires a protected implementation of the hash function SHA-512, whereas this is not required with implementations of the hash function SHA-256 used with ECDSA256 and ECDSA25519.

7.7. Cross-Algorithm and Cross-Protocol Attacks

The keypair used in this specification can be self-generated and the public key does not need to be exchanged, e.g., through certificates, with a third party before it is used.

New keypairs can be formed for new registration as the node desires. On the other hand, it is safer to allocate a keypair that is used only for the address protection and only for one instantiation of the signature scheme (which includes choice of elliptic curve domain parameters, used hash function, and applicable representation conventions).

The same private key MUST NOT be reused with more than one instantiation of the signature scheme in this specification. The same private key MUST NOT be used for anything other than computing NDPSO signatures per this specification.

ECDSA shall be used strictly as specified in [FIPS186-4]. In particular, each signing operation of ECDSA MUST use randomly generated ephemeral private keys and MUST NOT reuse these ephemeral private keys k accross signing operations. This precludes the use of deterministic ECDSA without a random input for determination of k , which is deemed dangerous for the intended applications this document aims to serve.

7.8. Public Key Validation

Public keys contained in the CIP0 field (which are used for signature verification) shall be verified to be correctly formed, by checking that this public key is indeed a point of the elliptic curve indicated by the Crypto-Type and that this point does have the proper order.

For points used with the signature scheme Ed25519, one MUST check that this point is not a point in the small subgroup (see Appendix B.1 of [CURVE-REPR]); for points used with the signature scheme ECDSA (i.e., both ECDSA256 and ECDSA25519), one MUST check that the point has the same order as the base point of the curve in question. This is commonly called full public key validation (again, see Appendix B.1 of [CURVE-REPR]).

7.9. Correlating Registrations

The ROVR field in the EARO introduced in [RFC8505] extends the EUI-64 field of the ARO defined in [RFC6775]. One of the drawbacks of using an EUI-64 as ROVR is that an attacker that is aware of the registrations can correlate traffic for a same 6LN across multiple addresses. Section 3 of [RFC8505] indicates that the ROVR and the address being registered are decoupled. A 6LN may use a same ROVR for multiple registrations or a different ROVR per registration, and the IID must not derive from the ROVR. In theory different 6LNs could use a same ROVR as long as they do not attempt to register the same address.

The Modifier used in the computation of the Crypto-ID enables a 6LN to build different Crypto-IDs for different addresses with a same keypair. Using that facility improves the privacy of the 6LN as the expense of storage in the 6LR, which will need to store multiple CIP0s that contain the same public key. Note that if the attacker is the 6LR, then the Modifier alone does not provide a protection, and the 6LN would need to use different keys and MAC addresses in an attempt to obfuscate its multiple ownership.

8. IANA considerations

8.1. CGA Message Type

This document defines a new 128-bit value of a Message Type tag under the CGA Message Type [RFC3972] name space: 0x8701 55c8 0cca dd32 6ab7 e415 f148 84d0.

8.2. Crypto-Type Subregistry

IANA is requested to create a new subregistry "Crypto-Type Subregistry" in the "Internet Control Message Protocol version 6 (ICMPv6) Parameters". The registry is indexed by an integer in the interval 0..255 and contains an Elliptic Curve, a Hash Function, a Signature Algorithm, Representation Conventions, Public key size, and Signature size, as shown in Table 1, which together specify a signature scheme (and which are fully specified in Appendix B).

The following Crypto-Type values are defined in this document:

Crypto-Type value	0 (ECDSA256)	1 (Ed25519)	2 (ECDSA25519)
Elliptic curve	NIST P-256 [FIPS186-4]	Curve25519 [RFC7748]	Curve25519 [RFC7748]
Hash function	SHA-256 [RFC6234]	SHA-512 [RFC6234]	SHA-256 [RFC6234]
Signature algorithm	ECDSA [FIPS186-4]	Ed25519 [RFC8032]	ECDSA [FIPS186-4]
Representation conventions	Weierstrass, (un)compressed, MSB/msb first, [RFC7518]	Edwards, compressed, LSB/lsb first, [RFC8037]	Weierstrass, (un)compressed, MSB/msb first, [CURVE-REPR]
Public key size	33/65 bytes (compressed/ uncompressed)	32 bytes (compressed)	33/65 bytes (compressed/ uncompressed)
Signature size	64 bytes	64 bytes	64 bytes
Defining specification	This_RFC	This_RFC	This_RFC

Table 1: Crypto-Types

New Crypto-Type values providing similar or better security may be defined in the future.

Assignment of new values for new Crypto-Type MUST be done through IANA with either "Specification Required" or "IESG Approval" as defined in BCP 26 [RFC8126].

8.3. IPv6 ND option types

This document registers two new ND option types under the subregistry "IPv6 Neighbor Discovery Option Formats":

Option Name	Suggested Value	Reference
NDP Signature Option (NDPSO)	38	This document
Crypto-ID Parameters Option (CIPO)	39	This document

Table 2: New ND options

8.4. New 6LoWPAN Capability Bit

IANA is requested to make additions to the Subregistry for "6LoWPAN Capability Bits" created for [RFC7400] as follows:

Capability Bit	Description	Document
09	AP-ND Enabled (1 bit)	This_RFC

Table 3: New 6LoWPAN Capability Bit

9. Acknowledgments

Many thanks to Charlie Perkins for his in-depth review and constructive suggestions. The authors are also especially grateful to Robert Moskowitz and Benjamin Kaduk for their comments and discussions that led to many improvements. The authors wish to also thank Shwetha Bhandari for actively shepherding this document and Roman Danyliw, Alissa Cooper, Mirja Kuhlewind, Eric Vyncke, Vijay Gurbani, Al Morton, and Adam Montville for their constructive reviews during the IESG process. Finally Many thanks to our INT area ADs, Suresh Krishnan and then Erik Kline, who supported us along the whole process.

10. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

- [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>>.
- [RFC6234] Eastlake 3rd, D. and T. Hansen, "US Secure Hash Algorithms (SHA and SHA-based HMAC and HKDF)", RFC 6234, DOI 10.17487/RFC6234, May 2011, <<https://www.rfc-editor.org/info/rfc6234>>.
- [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>>.
- [RFC7400] Bormann, C., "6LoWPAN-GHC: Generic Header Compression for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 7400, DOI 10.17487/RFC7400, November 2014, <<https://www.rfc-editor.org/info/rfc7400>>.
- [RFC7748] Langley, A., Hamburg, M., and S. Turner, "Elliptic Curves for Security", RFC 7748, DOI 10.17487/RFC7748, January 2016, <<https://www.rfc-editor.org/info/rfc7748>>.
- [RFC8032] Josefsson, S. and I. Liusvaara, "Edwards-Curve Digital Signature Algorithm (EdDSA)", RFC 8032, DOI 10.17487/RFC8032, January 2017, <<https://www.rfc-editor.org/info/rfc8032>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8505] Thubert, P., Ed., Nordmark, E., Chakrabarti, S., and C. Perkins, "Registration Extensions for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Neighbor Discovery", RFC 8505, DOI 10.17487/RFC8505, November 2018, <<https://www.rfc-editor.org/info/rfc8505>>.
- [FIPS186-4] FIPS 186-4, "Digital Signature Standard (DSS), Federal Information Processing Standards Publication 186-4", US Department of Commerce/National Institute of Standards and Technology , July 2013.
- [SEC1] SEC1, "SEC 1: Elliptic Curve Cryptography, Version 2.0", Standards for Efficient Cryptography , June 2009.

11. Informative references

- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, DOI 10.17487/RFC3972, March 2005, <<https://www.rfc-editor.org/info/rfc3972>>.
- [BCP 106] Eastlake 3rd, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", BCP 106, RFC 4086, DOI 10.17487/RFC4086, June 2005, <<https://www.rfc-editor.org/info/rfc4086>>.
- [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>>.
- [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>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<https://www.rfc-editor.org/info/rfc4944>>.
- [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, <<https://www.rfc-editor.org/info/rfc6282>>.
- [RFC7039] Wu, J., Bi, J., Bagnulo, M., Baker, F., and C. Vogt, Ed., "Source Address Validation Improvement (SAVI) Framework", RFC 7039, DOI 10.17487/RFC7039, October 2013, <<https://www.rfc-editor.org/info/rfc7039>>.
- [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, <<https://www.rfc-editor.org/info/rfc7217>>.

- [RFC7518] Jones, M., "JSON Web Algorithms (JWA)", RFC 7518, DOI 10.17487/RFC7518, May 2015, <<https://www.rfc-editor.org/info/rfc7518>>.
- [BCP 201] Housley, R., "Guidelines for Cryptographic Algorithm Agility and Selecting Mandatory-to-Implement Algorithms", BCP 201, RFC 7696, DOI 10.17487/RFC7696, November 2015, <<https://www.rfc-editor.org/info/rfc7696>>.
- [RFC8037] Liusvaara, I., "CFRG Elliptic Curve Diffie-Hellman (ECDH) and Signatures in JSON Object Signing and Encryption (JOSE)", RFC 8037, DOI 10.17487/RFC8037, January 2017, <<https://www.rfc-editor.org/info/rfc8037>>.
- [RFC8064] Gont, F., Cooper, A., Thaler, D., and W. Liu, "Recommendation on Stable IPv6 Interface Identifiers", RFC 8064, DOI 10.17487/RFC8064, February 2017, <<https://www.rfc-editor.org/info/rfc8064>>.
- [RFC8065] Thaler, D., "Privacy Considerations for IPv6 Adaptation-Layer Mechanisms", RFC 8065, DOI 10.17487/RFC8065, February 2017, <<https://www.rfc-editor.org/info/rfc8065>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [BACKBONE-ROUTER]
Thubert, P., Perkins, C., and E. Levy-Abegnoli, "IPv6 Backbone Router", Work in Progress, Internet-Draft, draft-ietf-6lo-backbone-router-20, 23 March 2020, <<https://tools.ietf.org/html/draft-ietf-6lo-backbone-router-20>>.
- [CURVE-REPR]
Struik, R., "Alternative Elliptic Curve Representations", Work in Progress, Internet-Draft, draft-ietf-lwig-curve-representations-09, 9 March 2020, <<https://tools.ietf.org/html/draft-ietf-lwig-curve-representations-09>>.
- [breaking-ed25519]
Samwel, N., Batina, L., Bertoni, G., Daemen, J., and R. Susella, "Breaking Ed25519 in WolfSSL", Cryptographers' Track at the RSA Conference , 2018, <https://link.springer.com/chapter/10.1007/978-3-319-76953-0_1>.

Appendix A. Requirements Addressed in this Document

In this section we state requirements of a secure neighbor discovery protocol for low-power and lossy networks.

- * The protocol MUST be based on the Neighbor Discovery Optimization for Low-power and Lossy Networks protocol defined in [RFC6775]. RFC6775 utilizes optimizations such as host-initiated interactions for sleeping resource-constrained hosts and elimination of multicast address resolution.
- * New options to be added to Neighbor Solicitation messages MUST lead to small packet sizes, especially compared with existing protocols such as SEcure Neighbor Discovery (SEND). Smaller packet sizes facilitate low-power transmission by resource-constrained nodes on lossy links.
- * The support for this registration mechanism SHOULD be extensible to more LLN links than IEEE 802.15.4 only. Support for at least the LLN links for which a 6lo "IPv6 over foo" specification exists, as well as Low-Power Wi-Fi SHOULD be possible.
- * 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.
- * The Address Registration Option used in the ND registration SHOULD be extended to carry the relevant forms of Unique Interface Identifier.
- * The Neighbor Discovery should specify the formation of a site-local address that follows the security recommendations from [RFC7217].

Appendix B. Representation Conventions

B.1. Signature Schemes

The signature scheme ECDSA256 corresponding to Crypto-Type 0 is ECDSA, as specified in [FIPS186-4], instantiated with the NIST prime curve P-256, as specified in Appendix B of [FIPS186-4], and the hash function SHA-256, as specified in [RFC6234], where points of this NIST curve are represented as points of a short-Weierstrass curve (see [FIPS186-4]) and are encoded as octet strings in most-significant-bit first (msb) and most-significant-byte first (MSB) order. The signature itself consists of two integers (r and s), which are each encoded as fixed-size octet strings in most-significant-bit first and most-significant-byte first order. For details on ECDSA, see [FIPS186-4]; for details on the encoding of public keys, see Appendix B.3; for details on the signature encoding, see Appendix B.2.

The signature scheme Ed25519 corresponding to Crypto-Type 1 is EdDSA, as specified in [RFC8032], instantiated with the Montgomery curve Curve25519, as specified in [RFC7748], and the hash function SHA-512, as specified in [RFC6234], where points of this Montgomery curve are represented as points of the corresponding twisted Edwards curve Edwards25519 (see Appendix B.4) and are encoded as octet strings in least-significant-bit first (lsb) and least-significant-byte first (LSB) order. The signature itself consists of a bit string that encodes a point of this twisted Edwards curve, in compressed format, and an integer encoded in least-significant-bit first and least-significant-byte first order. For details on EdDSA, the encoding of public keys and that of signatures, see the specification of pure Ed25519 in [RFC8032].

The signature scheme ECDSA25519 corresponding to Crypto-Type 2 is ECDSA, as specified in [FIPS186-4], instantiated with the Montgomery curve Curve25519, as specified in [RFC7748], and the hash function SHA-256, as specified in [RFC6234], where points of this Montgomery curve are represented as points of the corresponding short-Weierstrass curve Wei25519 (see Appendix B.4) and are encoded as octet strings in most-significant-bit first and most-significant-byte first order. The signature itself consists of a bit string that encodes two integers, each encoded as fixed-size octet strings in most-significant-bit first and most-significant-byte first order. For details on ECDSA, see [FIPS186-4]; for details on the encoding of public keys, see Appendix B.3; for details on the signature encoding, see Appendix B.2

B.2. Representation of ECDSA Signatures

With ECDSA, each signature is an ordered pair (r, s) of integers [FIPS186-4], where each integer is represented as a 32-octet string according to the Field Element to Octet String conversion rules in [SEC1] and where the ordered pair of integers is represented as the rightconcatenation of these representation values (thereby resulting in a 64-octet string). The inverse operation checks that the signature is a 64-octet string and represents the left-side and right-side halves of this string (each a 32-octet string) as the integers r and s , respectively, using the Octet String to Field Element conversion rules in [SEC1].

B.3. Representation of Public Keys Used with ECDSA

ECDSA is specified to be used with elliptic curves in short-Weierstrass form. Each point of such a curve is represented as an octet string using the Elliptic Curve Point to Octet String conversion rules in [SEC1], where point compression may be enabled (which is indicated by the leftmost octet of this representation). The inverse operation converts an octet string to a point of this curve using the Octet String to Elliptic Curve Point conversion rules in [SEC1], whereby the point is rejected if this is the so-called point at infinity. (This is the case if the input to this inverse operation is an octet string of length 1.)

B.4. Alternative Representations of Curve25519

The elliptic curve Curve25519, as specified in [RFC7748], is a so-called Montgomery curve. Each point of this curve can also be represented as a point of a twisted Edwards curve or as a point of an elliptic curve in short-Weierstrass form, via a coordinate transformation (a so-called isomorphic mapping). The parameters of the Montgomery curve and the corresponding isomorphic curves in twisted Edwards curve and short-Weierstrass form are as indicated below. Here, the domain parameters of the Montgomery curve Curve25519 and of the twisted Edwards curve Edwards25519 are as specified in [RFC7748]; the domain parameters of the elliptic curve Wei25519 in short-Weierstrass curve comply with Section 6.1.1 of [FIPS186-4]. For further details on these curves and on the coordinate transformations referenced above, see [CURVE-REPR].

General parameters (for all curve models):

```
p  2^{255}-19
    (=0xffffffff ffffffff ffffffff ffffffff ffffffff ffffffff
    ffffffffed)
h  8
n
723700557733226221397318656304299424085711635937990760600195093828
5454250989
(=2^{252} + 0x14def9de a2f79cd6 5812631a 5cf5d3ed)
```

Montgomery curve-specific parameters (for Curve25519):

```
A  486662
B  1
Gu 9 (=0x9)
```

Gv

```
147816194475895447910205935684099868872646061346164752889648818377
55586237401
(=0x20ae19a1 b8a086b4 e01edd2c 7748d14c 923d4d7e 6d7c61b2 29e9c5a2
7eced3d9)
```

Twisted Edwards curve-specific parameters (for Edwards25519):

a -1 (-0x01)

d -121665/121666
(=3709570593466943934313808350875456518954211387984321901638878553
3085940283555)
(=0x52036cee 2b6ffe73 8cc74079 7779e898 00700a4d 4141d8ab 75eb4dca
135978a3)

Gx

```
151122213495354007725011514095885315114540126930418572060461132839
49847762202
(=0x216936d3 cd6e53fe c0a4e231 fdd6dc5c 692cc760 9525a7b2 c9562d60
8f25d51a)
```

Gy 4/5

```
(=4631683569492647816942839400347516314130799386625622561578303360
3165251855960)
(=0x66666666 66666666 66666666 66666666 66666666 66666666 66666666
66666658)
```

Weierstrass curve-specific parameters (for Wei25519):

a

```
192986815395526992372618308347813179755449974442734273399095973345
73241639236
(=0x2aaaaaaaa aaaaaaaaa aaaaaaaaa aaaaaaaaa aaaaaaaaa aaaaaaa98
4914a144)
```

b

```
557517466698189089076452890782571408182411037279010123152944008379
56729358436
(=0x7b425ed0 97b425ed 097b425e d097b425 ed097b42 5ed097b4 260b5e9c
7710c864)
```

GX

```
192986815395526992372618308347813179755449974442734273399095973346
52188435546
(=0x2aaaaaaaa aaaaaaaaa aaaaaaaaa aaaaaaaaa aaaaaaaaa aaaaaaaaa
aaad245a)
```

GY

```
147816194475895447910205935684099868872646061346164752889648818377
55586237401
```

(=0x20ae19a1 b8a086b4 e01edd2c 7748d14c 923d4d7e 6d7c61b2 29e9c5a2
7eced3d9)

Authors' Addresses

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

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

Behcet Sarikaya

Email: sarikaya@ieee.org

Mohit Sethi
Ericsson
FI-02420 Jorvas
Finland

Email: mohit@piuha.net

Rene Struik
Struik Security Consultancy

Email: rstruik.ext@gmail.com

610
Internet-Draft
Updates: 6775, 8505 (if approved)
Intended status: Standards Track
Expires: 24 September 2020

P. Thubert, Ed.
Cisco Systems
C.E. Perkins
Blue Meadow Networking
E. Levy-Abegnoli
Cisco Systems
23 March 2020

IPv6 Backbone Router
draft-ietf-610-backbone-router-20

Abstract

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

Status of This Memo

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

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

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

This Internet-Draft will expire on 24 September 2020.

Copyright Notice

Copyright (c) 2020 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. Terminology	5
2.1. BCP 14	5
2.2. New Terms	5
2.3. Abbreviations	6
2.4. References	7
3. Overview	7
3.1. Updating RFC 6775 and RFC 8505	10
3.2. Access Link	11
3.3. Route-Over Mesh	13
3.4. The Binding Table	14
3.5. Primary and Secondary 6BBRs	15
3.6. Using Optimistic DAD	16
4. Multi-Link Subnet Considerations	17
5. Optional 6LBR serving the Multi-Link Subnet	17
6. Using IPv6 ND Over the Backbone Link	18
7. Routing Proxy Operations	20
8. Bridging Proxy Operations	21
9. Creating and Maintaining a Binding	22
9.1. Operations on a Binding in Tentative State	23
9.2. Operations on a Binding in Reachable State	24
9.3. Operations on a Binding in Stale State	25
10. Registering Node Considerations	26
11. Security Considerations	27
12. Protocol Constants	30
13. IANA Considerations	30
14. Acknowledgments	30
15. Normative References	30
16. Informative References	32
Appendix A. Possible Future Extensions	34
Appendix B. Applicability and Requirements Served	35
Authors' Addresses	37

1. Introduction

IEEE STD. 802.1 [IEEEstd8021] Ethernet Bridging provides an efficient and reliable broadcast service for wired networks; applications and protocols have been built that heavily depend on that feature for their core operation. Unfortunately, Low-Power Lossy Networks (LLNs) and local wireless networks generally do not provide the broadcast capabilities of Ethernet Bridging in an economical fashion.

As a result, protocols designed for bridged networks that rely on multicast and broadcast often exhibit disappointing behaviours when employed unmodified on a local wireless medium (see [I-D.ietf-mboned-ieee802-mcast-problems]).

Wi-Fi [IEEEstd80211] Access Points (APs) deployed in an Extended Service Set (ESS) act as Ethernet Bridges [IEEEstd8021], with the property that the bridging state is established at the time of association. This ensures connectivity to the end node (the Wi-Fi STA) and protects the wireless medium against broadcast-intensive Transparent Bridging reactive Lookups. In other words, the association process is used to register the MAC Address of the STA to the AP. The AP subsequently proxies the bridging operation and does not need to forward the broadcast Lookups over the radio.

In the same way as Transparent Bridging, IPv6 [RFC8200] Neighbor Discovery [RFC4861] [RFC4862] Protocol (IPv6 ND) is a reactive protocol, based on multicast transmissions to locate an on-link correspondent and ensure the uniqueness of an IPv6 address. The mechanism for Duplicate Address Detection (DAD) [RFC4862] was designed for the efficient broadcast operation of Ethernet Bridging. Since broadcast can be unreliable over wireless media, DAD often fails to discover duplications [I-D.yourtchenko-6man-dad-issues]. In practice, the fact that IPv6 addresses very rarely conflict is mostly attributable to the entropy of the 64-bit Interface IDs as opposed to the successful operation of the IPv6 ND duplicate address detection and resolution mechanisms.

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

Because IPv6 ND messages sent to the SNMA group are broadcast at the radio MAC Layer, wireless nodes that do not belong to the SNMA group still have to keep their radio turned on to listen to multicast NS messages, which is a waste of energy for them. In order to reduce their power consumption, certain battery-operated devices such as IoT sensors and smartphones ignore some of the broadcasts, making IPv6 ND operations even less reliable.

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

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

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

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

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

Knowledge of which address to proxy for can be obtained by snooping the IPV6 ND protocol (see [I-D.bi-savi-wlan]), but it has been found to be unreliable. An IPv6 address may not be discovered immediately due to a packet loss, or if a "silent" node is not currently using one of its addresses. A change of state (e.g., due to movement) may be missed or misordered, leading to unreliable connectivity and incomplete knowledge of the state of the network.

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

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

2. Terminology

2.1. BCP 14

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2.2. New Terms

This document introduces the following terminology:

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

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

Routing Proxy: A Routing Proxy provides IPv6 ND proxy functions and enables the MLSN operation over federated links that may not be compatible for bridging. The Routing Proxy advertises its own MAC Address as the Target Link Layer Address (TLLA) in the proxied NAs

over the Backbone, and routes at the Network Layer between the federated links.

Bridging Proxy: A Bridging Proxy provides IPv6 ND proxy functions while preserving forwarding continuity at the MAC Layer. In that case, the MAC Address and the mobility of the Registering Node is visible across the bridged Backbone. The Bridging Proxy advertises the MAC Address of the Registering Node as the TLLA in the proxied NAs over the Backbone, and proxies ND for all unicast addresses including Link-Local Addresses. Instead of replying on behalf of the Registering Node, a Bridging Proxy will preferably forward the NS Lookup and NUD messages that target the Registered Address to the Registering Node as unicast frames and let it respond in its own.

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

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

2.3. Abbreviations

This document uses the following abbreviations:

6BBR: 6LoWPAN Backbone Router
6LBR: 6LoWPAN Border Router
6LN: 6LoWPAN Node
6LR: 6LoWPAN Router
ARO: Address Registration Option
DAC: Duplicate Address Confirmation
DAD: Duplicate Address Detection
DAR: Duplicate Address Request
EARO: Extended Address Registration Option
EDAC: Extended Duplicate Address Confirmation
EDAR: Extended Duplicate Address Request
DODAG: Destination-Oriented Directed Acyclic Graph
ID: Identifier
LLN: Low-Power and Lossy Network
NA: Neighbor Advertisement
MAC: Medium Access Control
NCE: Neighbor Cache Entry
ND: Neighbor Discovery
NDP: Neighbor Discovery Protocol
NS: Neighbor Solicitation

NS(DAD): NDP NS message used for the purpose of duplication avoidance (multicast)
NS(Lookup): NDP NS message used for the purpose of address resolution (multicast)
NS(NUD): NDP NS message used for the purpose of unreachability detection (unicast)
NUD: Neighbor Unreachability Detection
ROVR: Registration Ownership Verifier
RPL: IPv6 Routing Protocol for LLNs
RA: Router Advertisement
RS: Router Solicitation
SNMA: Solicited-Node Multicast Address
LLA: Link Layer Address (aka MAC address)
SLLA: Source Link Layer Address
TLLA: Target Link Layer Address
TID: Transaction ID

2.4. References

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

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

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

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

6LoWPAN ND: Neighbor Discovery Optimization for Low-Power and Lossy Networks [RFC6775], "Registration Extensions for 6LoWPAN Neighbor Discovery" [RFC8505], and "Address Protected Neighbor Discovery for Low-power and Lossy Networks" [I-D.ietf-6lo-ap-nd].

3. Overview

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

Figure 1 illustrates a backbone link that federates a collection of LLNs as a single IPv6 Subnet, with a number of 6BBRs providing proxy-ND services to their attached LLNs.

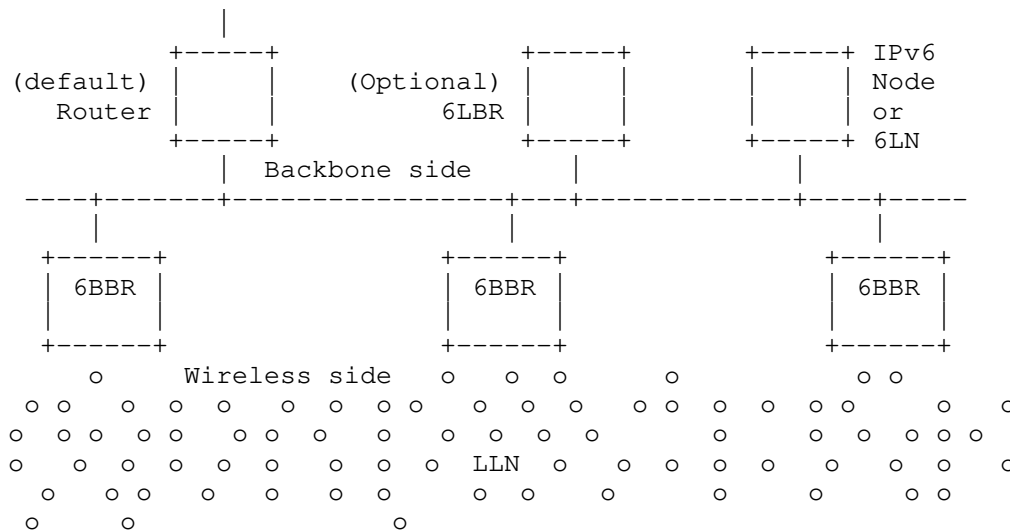


Figure 1: Backbone Link and Backbone Routers

The LLN may be a hub-and-spoke access link such as (Low-Power) IEEE STD. 802.11 (Wi-Fi) [IEEEstd80211] and IEEE STD. 802.15.1 (Bluetooth) [IEEEstd802151], or a Mesh-Under or a Route-Over network [RFC8505]. The proxy state can be distributed across multiple 6BBRs attached to the same Backbone.

The main features of a 6BBR are as follows:

- * Multi-Link-subnet functions (provided by the 6BBR on the backbone) performed on behalf of Registered Nodes, and
- * Routing registrar services that reduce multicast within the LLN:
 - Binding Table management
 - failover, e.g., due to mobility

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

Unless otherwise configured, a 6BBR does the following:

- * Create a new entry in a Binding Table for a new Registered Address and ensure that the Address is not duplicated over the Backbone.
- * Advertise a Registered Address over the Backbone using an NA message, either unsolicited or as a response to a NS message. This includes joining the multicast group associated to the SNMA derived from the Registered Address as specified in section 7.2.1. of [RFC4861] over the Backbone.
- * The 6BBR MAY respond immediately as a Proxy in lieu of the Registering Node, e.g., if the Registering Node has a sleeping cycle that the 6BBR does not want to interrupt, or if the 6BBR has a recent state that is deemed fresh enough to permit the proxied response. It is preferred, though, that the 6BBR checks whether the Registering Node is still responsive on the Registered Address. To that effect:
 - as a Bridging Proxy:
the 6BBR forwards the multicast DAD and Address Lookup messages as a unicast MAC-Layer frames to the MAC address of the Registering Node that matches the Target in the ND message, and forwards as is the unicast Neighbor Unreachability Detection (NUD) messages, so as to let the Registering Node answer with the ND Message and options that it sees fit;
 - as a Routing Proxy:
the 6BBR checks the liveliness of the Registering Node, e.g., using a NUD verification, before answering on its behalf.
- * Deliver packets arriving from the LLN, using Neighbor Solicitation messages to look up the destination over the Backbone.
- * Forward or bridge packets between the LLN and the Backbone.
- * Verify liveness for a registration, when needed.

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

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

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

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

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

- * The EARO [RFC8505] is used in the IPv6 ND exchanges over the Backbone between the 6BBRs to help distinguish duplication from movement. Extended Duplicate Address Messages (EDAR and EDAC) may also be used to communicate with a 6LBR, if one is present. Address duplication is detected using the ROVR field. Conflicting registrations to different 6BBRs for the same Registered Address are resolved using the TID field which forms an order of registrations.
- * The Link Layer Address (LLA) that the 6BBR advertises for the Registered Address on behalf of the Registered Node over the Backbone can belong to the Registering Node; in that case, the 6BBR (acting as a Bridging Proxy (see Section 8)) bridges the unicast packets. Alternatively, the LLA can be that of the 6BBR on the Backbone interface, in which case the 6BBR (acting as a Routing Proxy (see Section 7)) receives the unicast packets at Layer 3 and routes over.

3.1. Updating RFC 6775 and RFC 8505

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

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

This specification allows to deploy a 6LBR on the backbone where EDAR and EDAC messages coexist with classical ND. It also adds the capability to insert IPv6 ND options in the EDAR and EDAC messages. A 6BBR acting as a 6LR for the Registered Address can insert an SLLAO

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

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

3.2. Access Link

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

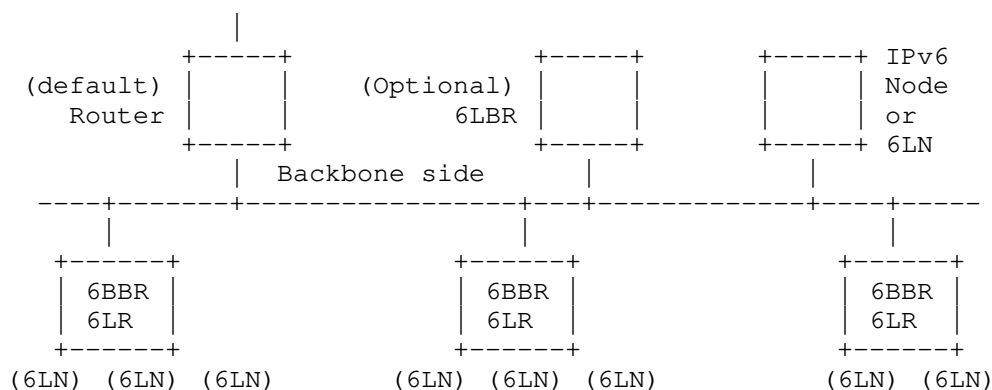


Figure 2: Access Link Use case

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

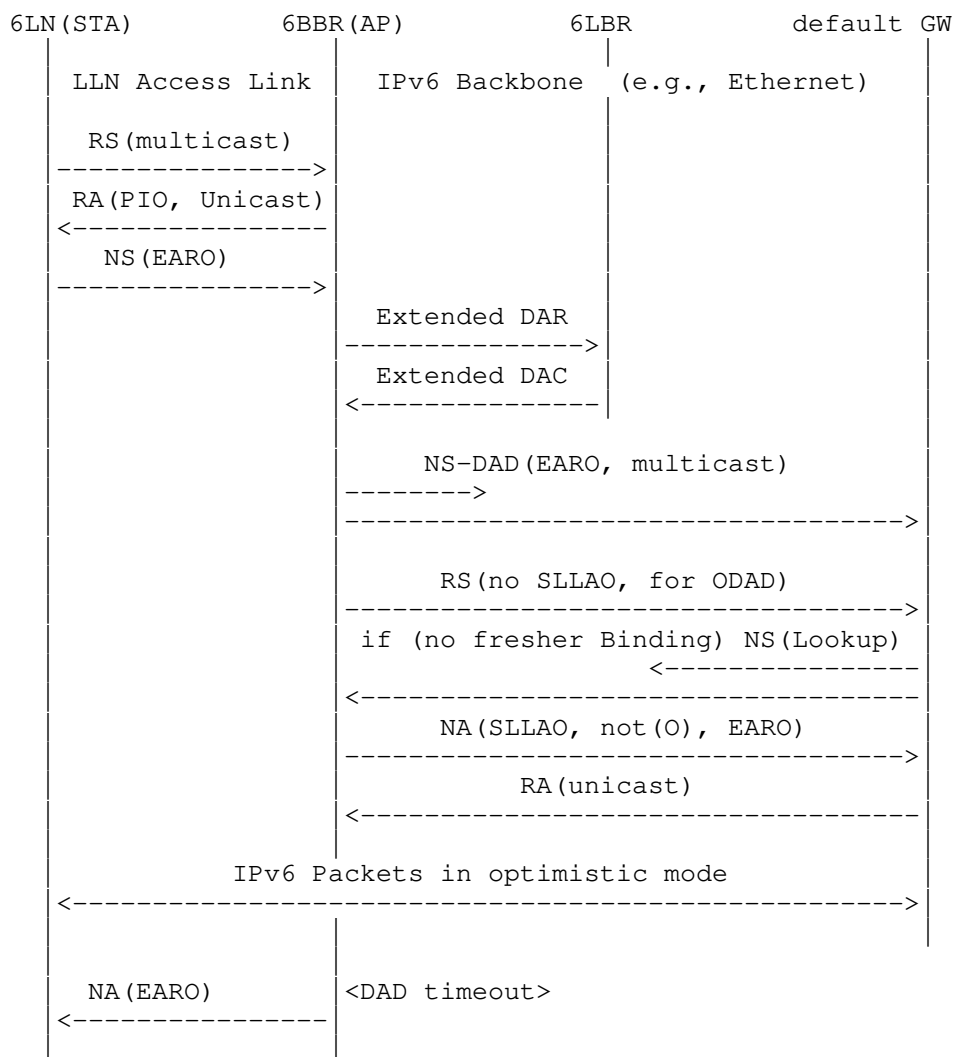


Figure 3: Initial Registration Flow to a 6BBR acting as Routing Proxy

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

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

3.3. Route-Over Mesh

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

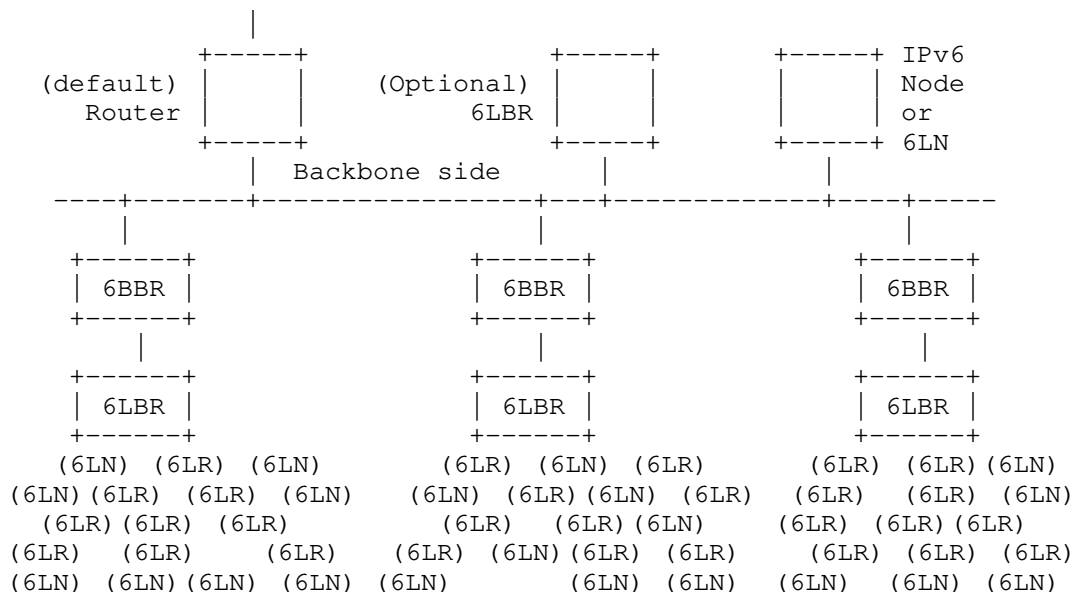


Figure 4: Route-Over Mesh Use case

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

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

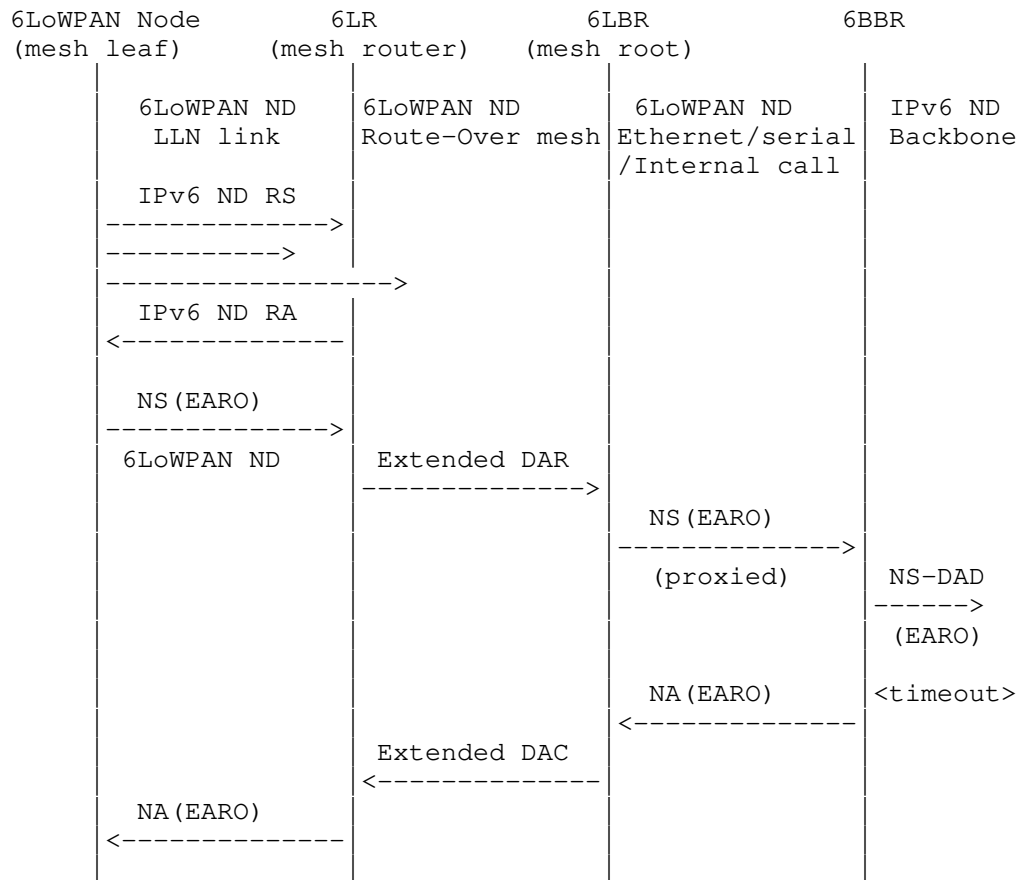


Figure 5: Initial Registration Flow over Route-Over Mesh

As a non-normative example of a Route-Over Mesh, the 6TiSCH architecture [I-D.ietf-6tisch-architecture] suggests using the RPL [RFC6550] routing protocol and collocating the RPL root with a 6LBR that serves the LLN. The 6LBR is also either collocated with or directly connected to the 6BBR over an IPv6 Link.

3.4. The Binding Table

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

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

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

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

Over the LLN, Binding Table management is as follows:

- * De-registrations (newer TID, same ROVR, null Lifetime) are accepted with a status of 4 ("Removed"); the entry is deleted;
- * Newer registrations (newer TID, same ROVR, non-null Lifetime) are accepted with a status of 0 (Success); the Binding is updated with the new TID, the Registration Lifetime and the Registering Node; in Tentative state the EDAC response is held and may be overwritten; in other states the Registration Lifetime timer is restarted and the entry is placed in Reachable state.
- * Identical registrations (same TID, same ROVR) from the same Registering Node are accepted with a status of 0 (Success). In Tentative state, the response is held and may be overwritten, but the response is eventually produced, carrying the result of the DAD process;
- * Older registrations (older TID, same ROVR) from the same Registering Node are discarded;
- * Identical and older registrations (not-newer TID, same ROVR) from a different Registering Node are rejected with a status of 3 (Moved); this may be rate limited to avoid undue interference;
- * Any registration for the same address but with a different ROVR is rejected with a status of 1 (Duplicate).

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

3.5. Primary and Secondary 6BBRs

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

another 6BBR (same Registered address, same TID, same ROVR) are silently ignored but for the purpose of selecting the primary 6BBR for that registration.

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

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

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

3.6. Using Optimistic DAD

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

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

Following ODAD, the router may then send a unicast RA to the Registered Address, and it may resolve that Address using an NS(Lookup) message. In response, the 6BBR sends an NA with an EARO and the Override flag [RFC4861] that is not set. The router can then determine the freshest EARO in case of conflicting NA(EARO) messages, using the method described in section 5.2.1 of [RFC8505]. If the NA(EARO) is the freshest answer, the default router creates a Binding

with the SLIAO of the 6BBR (in Routing Proxy mode) or that of the Registering Node (in Bridging Proxy mode) so that traffic from/to the Registered Address can flow immediately.

4. Multi-Link Subnet Considerations

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

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

5. Optional 6LBR serving the Multi-Link Subnet

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

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

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

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

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

6. Using IPv6 ND Over the Backbone Link

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

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

IPv6 ND operates as follows on the backbone:

- * Section 7.2.8 of [RFC4861] specifies that an NA message generated as a proxy does not have the Override flag set in order to ensure that if the real owner is present on the link, its own NA will take precedence, and that this NA does not update the NCE for the real owner if one exists.
- * A node that receives multiple NA messages updates an existing NCE only if the Override flag is set; otherwise the node will probe the cached address.
- * When an NS(DAD) is received for a tentative address, which means that two nodes form the same address at nearly the same time, section 5.4.3 of [RFC4862] cannot detect which node first claimed the address and the address is abandoned.
- * In any case, [RFC4862] indicates that a node never responds to a Neighbor Solicitation for a tentative address.

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

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

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

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

- * an NS(DAD) that indicates a duplication can still not be asserted for first come, but the situation can be avoided using a 6LBR on the backbone that will serialize the order of appearance of the address and ensure first-come/first-serve.
- * an NS or an NA that denotes an older registration for the same Registered Node is not interpreted as a duplication as specified in section 5.4.3 and 5.4.4 of [RFC4862], respectively.

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

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

In any state, the 6BBR acts as follows:

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

This behavior is specified in more detail in Section 9.

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

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

7. Routing Proxy Operations

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

In this mode, it is not required that the MAC addresses of the 6LNs are visible at Layer 2 over the Backbone. It is thus useful when the messaging over the Backbone that is associated to wireless mobility becomes expensive, e.g., when the Layer 2 topology is virtualized over a wide area IP underlay.

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

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

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

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

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

8. Bridging Proxy Operations

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

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

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

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

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

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

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

9. Creating and Maintaining a Binding

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

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

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

- * Acting as Bridging Proxy the 6LR MUST proxy ND over the backbone for registered Link-Local Addresses.
- * EDAR and EDAC messages SHOULD carry a SLLAO and a TLLAO, respectively.
- * An EDAC message with a status of 9 (6LBR Registry Saturated) is assimilated as a status of 0 if a following DAD process protects the address against duplication.

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

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

If a registration is received for an existing Binding with a non-null Registration Lifetime and the registration is fresher (same ROVR, fresher TID), then the Binding is updated, with the new Registration Lifetime, TID, and possibly Registering Node. In Tentative state

(see Section 9.1), the current DAD operation continues unaltered. In other states (see Section 9.2 and Section 9.3), the Binding is placed in Reachable state for the Registration Lifetime, and the 6BBR returns an NA(EARO) to the Registering Node with a status of 0 (Success).

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

If a registration is received for an existing Binding and a registration Lifetime of zero, then the Binding is removed, and the 6BBR returns an NA(EARO) back to the Registering Node with a status of 0 (Success). An implementation of a Routing Proxy that removes a binding MUST remove the associated Host route pointing on the registering node.

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

9.1. Operations on a Binding in Tentative State

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

For a Binding in Tentative state:

- * The Binding MUST be removed if an NA message is received over the Backbone for the Registered Address with no EARO, or containing an EARO that indicates an existing registration owned by a different Registering Node (different ROVR). In that case, an NA is sent back to the Registering Node with a status of 1 (Duplicate) to indicate that the binding has been rejected. This behavior might be overridden by policy, in particular if the registration is trusted, e.g., based on the validation of the ROVR field (see [I-D.ietf-6lo-ap-nd]).

- * The Binding MUST be removed if an NS(DAD) message is received over the Backbone for the Registered Address with no EARO, or containing an EARO with a different ROVR that indicates a tentative registration by a different Registering Node. In that case, an NA is sent back to the Registering Node with a status of 1 (Duplicate). This behavior might be overridden by policy, in particular if the registration is trusted, e.g., based on the validation of the ROVR field (see [I-D.ietf-6lo-ap-nd]).
- * The Binding MUST be removed if an NA or an NS(DAD) message is received over the Backbone for the Registered Address containing an EARO with a that indicates a fresher registration ([RFC8505]) for the same Registering Node (same ROVR). In that case, an NA MUST be sent back to the Registering Node with a status of 3 (Moved).
- * The Binding MUST be kept unchanged if an NA or an NS(DAD) message is received over the Backbone for the Registered Address containing an EARO with a that indicates an older registration ([RFC8505]) for the same Registering Node (same ROVR). The message is answered with an NA that carries an EARO with a status of 3 (Moved) and the Override flag not set. This behavior might be overridden by policy, in particular if the registration is not trusted.
- * Other NS(DAD) and NA messages from the Backbone are ignored.
- * NS(Lookup) and NS(NUD) messages SHOULD be optimistically answered with an NA message containing an EARO with a status of 0 and the Override flag not set (see Section 3.6). If optimistic DAD is disabled, then they SHOULD be queued to be answered when the Binding goes to Reachable state.

When the TENTATIVE_DURATION (Section 12) timer elapses, the Binding is placed in Reachable state for the Registration Lifetime, and the 6BBR returns an NA(EARO) to the Registering Node with a status of 0 (Success).

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

9.2. Operations on a Binding in Reachable State

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

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

For a Binding in Reachable state:

- * The Binding MUST be removed if an NA or an NS(DAD) message is received over the Backbone for the Registered Address containing an EARO that indicates a fresher registration ([RFC8505]) for the same Registered Node (i.e., same ROVR but fresher TID). A status of 4 (Removed) is returned in an asynchronous NA(EARO) to the Registering Node. Based on configuration, an implementation may delay this operation by a timer with a short setting, e.g., a few seconds to a minute, in order to allow for a parallel registration to reach this node, in which case the NA might be ignored.
- * NS(DAD) and NA messages containing an EARO that indicates a registration for the same Registered Node that is not as fresh as this binding MUST be answered with an NA message containing an EARO with a status of 3 (Moved).
- * An NS(DAD) with no EARO or with an EARO that indicates a duplicate registration (i.e., different ROVR) MUST be answered with an NA message containing an EARO with a status of 1 (Duplicate) and the Override flag not set, unless the received message is an NA that carries an EARO with a status of 1, in which case the node refrains from answering.
- * Other NS(DAD) and NA messages from the Backbone are ignored.
- * NS(Lookup) and NS(NUD) messages SHOULD be answered with an NA message containing an EARO with a status of 0 and the Override flag not set. The 6BBR MAY check whether the Registering Node is still available using a NUD procedure over the LLN prior to answering; this behaviour depends on the use case and is subject to configuration.

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

9.3. Operations on a Binding in Stale State

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

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

For a Binding in Stale state:

- * The Binding MUST be removed if an NA or an NS(DAD) message is received over the Backbone for the Registered Address containing no EARO or an EARO that indicates either a fresher registration for the same Registered Node or a duplicate registration. A status of 4 (Removed) MAY be returned in an asynchronous NA(EARO) to the Registering Node.
- * NS(DAD) and NA messages containing an EARO that indicates a registration for the same Registered Node that is not as fresh as this MUST be answered with an NA message containing an EARO with a status of 3 (Moved).
- * If the 6BBR receives an NS(Lookup) or an NS(NUD) message for the Registered Address, the 6BBR MUST attempt a NUD procedure as specified in [RFC7048] to the Registering Node, targeting the Registered Address, prior to answering. If the NUD procedure succeeds, the operation in Reachable state applies. If the NUD fails, the 6BBR refrains from answering.
- * Other NS(DAD) and NA messages from the Backbone are ignored.

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

10. Registering Node Considerations

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

The Registering Node may be the 6LN owning the IPv6 Address, or a 6LBR that performs the registration on its behalf in a Route-Over mesh.

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

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

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

11. Security Considerations

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

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

In particular, the LLN MAC is required to provide secure unicast to/from the Backbone Router and secure broadcast from the routers in a way that prevents tampering with or replaying the ND messages.

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

It results that the security threats that are detailed in section 11.1 of [RFC4861] fully apply to this specification as well. In very short:

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

- * Any node that can send a packet on the backbone can forge traffic and pretend it is issued from a address that it does not own, even if it did not claim the address using ND.
- * Any node that can send a packet on the backbone can present itself as a preferred router to intercept all traffic outgoing the subnet. It may even expose a prefix on the subnet as not-on-link and intercept all the traffic within the subnet.
- * If the rogue can receive a packet from the backbone it can also snoop all the intercepted traffic, be it by stealing an address or the role of a router.

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

Using address registration as the sole ND mechanism on a link and coupling it with [I-D.ietf-6lo-ap-nd] guarantees the ownership of a registered address within that link.

- * The protection is based on a proof-of-ownership encoded in the ROVR field and protects against address theft and impersonation by a 6LN, because the 6LR can challenge the Registered Node for a proof-of-ownership.
- * The protection extends to the full LLN in the case of an LLN Link, but does not extend over the backbone since the 6BBR cannot provide the proof-of-ownership when it defends the address.

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

If the classical ND is disabled on the backbone and the use of [I-D.ietf-6lo-ap-nd] and a 6LBR are mandated, the network will benefit from the following new advantages:

Zero-trust security for ND flows within the whole subnet: the increased security that [I-D.ietf-6lo-ap-nd] provides on the LLN will also apply to the backbone; it becomes impossible for an attached node to claim an address that belongs to another node using ND, and the network can filter packets that are not originated by the owner of the source address (SAVI), as long as that the routers are known and trusted.

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

Less IPv6 ND-related multicast on the backbone: DAD and NS(Lookup) become unicast queries to the 6LBR

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

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

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

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

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

12. Protocol Constants

This Specification uses the following constants:

TENTATIVE_DURATION: 800 milliseconds

STALE_DURATION: see below

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

13. IANA Considerations

This document has no request to IANA.

14. Acknowledgments

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

15. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.
- [RFC4429] Moore, N., "Optimistic Duplicate Address Detection (DAD) for IPv6", RFC 4429, DOI 10.17487/RFC4429, April 2006, <<https://www.rfc-editor.org/info/rfc4429>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861,

- DOI 10.17487/RFC4861, September 2007,
<<https://www.rfc-editor.org/info/rfc4861>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <<https://www.rfc-editor.org/info/rfc4862>>.
- [RFC6059] Krishnan, S. and G. Daley, "Simple Procedures for Detecting Network Attachment in IPv6", RFC 6059, DOI 10.17487/RFC6059, November 2010, <<https://www.rfc-editor.org/info/rfc6059>>.
- [RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <<https://www.rfc-editor.org/info/rfc6775>>.
- [RFC7048] Nordmark, E. and I. Gashinsky, "Neighbor Unreachability Detection Is Too Impatient", RFC 7048, DOI 10.17487/RFC7048, January 2014, <<https://www.rfc-editor.org/info/rfc7048>>.
- [RFC7559] Krishnan, S., Anipko, D., and D. Thaler, "Packet-Loss Resiliency for Router Solicitations", RFC 7559, DOI 10.17487/RFC7559, May 2015, <<https://www.rfc-editor.org/info/rfc7559>>.
- [RFC7772] Yourtchenko, A. and L. Colitti, "Reducing Energy Consumption of Router Advertisements", BCP 202, RFC 7772, DOI 10.17487/RFC7772, February 2016, <<https://www.rfc-editor.org/info/rfc7772>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.
- [RFC8201] McCann, J., Deering, S., Mogul, J., and R. Hinden, Ed., "Path MTU Discovery for IP version 6", STD 87, RFC 8201, DOI 10.17487/RFC8201, July 2017, <<https://www.rfc-editor.org/info/rfc8201>>.

- [RFC8505] Thubert, P., Ed., Nordmark, E., Chakrabarti, S., and C. Perkins, "Registration Extensions for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Neighbor Discovery", RFC 8505, DOI 10.17487/RFC8505, November 2018, <<https://www.rfc-editor.org/info/rfc8505>>.

16. Informative References

- [RFC4389] Thaler, D., Talwar, M., and C. Patel, "Neighbor Discovery Proxies (ND Proxy)", RFC 4389, DOI 10.17487/RFC4389, April 2006, <<https://www.rfc-editor.org/info/rfc4389>>.
- [RFC4903] Thaler, D., "Multi-Link Subnet Issues", RFC 4903, DOI 10.17487/RFC4903, June 2007, <<https://www.rfc-editor.org/info/rfc4903>>.
- [RFC5415] Calhoun, P., Ed., Montemurro, M., Ed., and D. Stanley, Ed., "Control And Provisioning of Wireless Access Points (CAPWAP) Protocol Specification", RFC 5415, DOI 10.17487/RFC5415, March 2009, <<https://www.rfc-editor.org/info/rfc5415>>.
- [RFC5568] Koodli, R., Ed., "Mobile IPv6 Fast Handovers", RFC 5568, DOI 10.17487/RFC5568, July 2009, <<https://www.rfc-editor.org/info/rfc5568>>.
- [RFC6606] Kim, E., Kaspar, D., Gomez, C., and C. Bormann, "Problem Statement and Requirements for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing", RFC 6606, DOI 10.17487/RFC6606, May 2012, <<https://www.rfc-editor.org/info/rfc6606>>.
- [RFC6275] Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, DOI 10.17487/RFC6275, July 2011, <<https://www.rfc-editor.org/info/rfc6275>>.
- [RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <<https://www.rfc-editor.org/info/rfc6550>>.
- [RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", RFC 6830, DOI 10.17487/RFC6830, January 2013, <<https://www.rfc-editor.org/info/rfc6830>>.

- [RFC8273] Brzozowski, J. and G. Van de Velde, "Unique IPv6 Prefix per Host", RFC 8273, DOI 10.17487/RFC8273, December 2017, <<https://www.rfc-editor.org/info/rfc8273>>.
- [I-D.yourtchenko-6man-dad-issues]
Yourtchenko, A. and E. Nordmark, "A survey of issues related to IPv6 Duplicate Address Detection", Work in Progress, Internet-Draft, draft-yourtchenko-6man-dad-issues-01, 3 March 2015, <<https://tools.ietf.org/html/draft-yourtchenko-6man-dad-issues-01>>.
- [I-D.nordmark-6man-dad-approaches]
Nordmark, E., "Possible approaches to make DAD more robust and/or efficient", Work in Progress, Internet-Draft, draft-nordmark-6man-dad-approaches-02, 19 October 2015, <<https://tools.ietf.org/html/draft-nordmark-6man-dad-approaches-02>>.
- [I-D.ietf-6man-rs-refresh]
Nordmark, E., Yourtchenko, A., and S. Krishnan, "IPv6 Neighbor Discovery Optional RS/RA Refresh", Work in Progress, Internet-Draft, draft-ietf-6man-rs-refresh-02, 31 October 2016, <<https://tools.ietf.org/html/draft-ietf-6man-rs-refresh-02>>.
- [I-D.ietf-6lo-ap-nd]
Thubert, P., Sarikaya, B., Sethi, M., and R. Struik, "Address Protected Neighbor Discovery for Low-power and Lossy Networks", Work in Progress, Internet-Draft, draft-ietf-6lo-ap-nd-20, 9 March 2020, <<https://tools.ietf.org/html/draft-ietf-6lo-ap-nd-20>>.
- [I-D.ietf-6tisch-architecture]
Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", Work in Progress, Internet-Draft, draft-ietf-6tisch-architecture-28, 29 October 2019, <<https://tools.ietf.org/html/draft-ietf-6tisch-architecture-28>>.
- [I-D.ietf-mboned-ieee802-mcast-problems]
Perkins, C., McBride, M., Stanley, D., Kumari, W., and J. Zuniga, "Multicast Considerations over IEEE 802 Wireless Media", Work in Progress, Internet-Draft, draft-ietf-mboned-ieee802-mcast-problems-11, 11 December 2019, <<https://tools.ietf.org/html/draft-ietf-mboned-ieee802-mcast-problems-11>>.

[I-D.bi-savi-wlan]

Bi, J., Wu, J., Wang, Y., and T. Lin, "A SAVI Solution for WLAN", Work in Progress, Internet-Draft, draft-bi-savi-wlan-18, 17 November 2019, <<https://tools.ietf.org/html/draft-bi-savi-wlan-18>>.

[I-D.thubert-6lo-unicast-lookup]

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

[IEEEStd8021]

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

[IEEEStd80211]

IEEE standard for Information Technology, "IEEE Standard for Information technology -- Telecommunications and information exchange between systems Local and metropolitan area networks-- Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications".

[IEEEStd802151]

IEEE standard for Information Technology, "IEEE Standard for Information Technology - Telecommunications and Information Exchange Between Systems - Local and Metropolitan Area Networks - Specific Requirements. - Part 15.1: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Wireless Personal Area Networks (WPANs)".

[IEEEStd802154]

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

Appendix A. Possible Future Extensions

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

By default the specification does not have a fine-grained trust model: all nodes that can authenticate to the LLN MAC or attach to the backbone are equally trusted. It would be desirable to provide a stronger authorization model, e.g., whereby nodes that associate their address with a proof-of-ownership [I-D.ietf-6lo-ap-nd] should be more trusted than nodes that do not. Such a trust model and related signaling could be added in the future to override the default operation and favor trusted nodes.

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

Appendix B. Applicability and Requirements Served

This document specifies proxy-ND functions that can be used to federate an IPv6 Backbone Link and multiple IPv6 LLNs into a single Multi-Link Subnet. The proxy-ND functions enable IPv6 ND services for Duplicate Address Detection (DAD) and Address Lookup that do not require broadcasts over the LLNs.

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

Each LLN in the subnet is attached to an IPv6 Backbone Router (6BBR). The Backbone Routers interconnect the LLNs and advertise the Addresses of the 6LNs over the Backbone Link using proxy-ND operations.

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

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

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

In the case of a Wi-Fi access link, a 6BBR may be collocated with the Access Point (AP), or with a Fabric Edge (FE) or a CAPWAP [RFC5415] Wireless LAN Controller (WLC). In those cases, the wireless client (STA) is the 6LN that makes use of [RFC8505] to register its IPv6 Address(es) to the 6BBR acting as Routing Registrar. The 6LBR can be centralized and either connected to the Backbone Link or reachable over IP. The 6BBR proxy-ND operations eliminate the need for wireless nodes to respond synchronously when a Lookup is performed for their IPv6 Addresses. This provides the function of a Sleep Proxy for ND [I-D.nordmark-6man-dad-approaches].

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

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

Scalability of the Multi-Link Subnet [RFC4903] requires avoidance of multicast/broadcast operations as much as possible even on the Backbone [I-D.ietf-mboned-ieee802-mcast-problems]. Although hosts can connect to the Backbone using IPv6 ND operations, multicast RAs can be saved by using [I-D.ietf-6man-rs-refresh], which also requires the support of [RFC7559].

Authors' Addresses

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

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

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

Email: charliep@computer.org

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

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

6Lo Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 14, 2018

C. Gomez
S. Darroudi
UPC/i2cat
T. Savolainen
Nokia
September 10, 2017

IPv6 Mesh over BLUETOOTH(R) Low Energy using IPSP
draft-ietf-6lo-blemesh-02

Abstract

RFC 7668 describes the adaptation of 6LoWPAN techniques to enable IPv6 over Bluetooth low energy networks that follow the star topology. However, recent Bluetooth specifications allow the formation of extended topologies as well. This document specifies the mechanisms needed to enable IPv6 over mesh networks composed of Bluetooth low energy links established by using the Bluetooth Internet Protocol Support Profile.

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 March 14, 2018.

Copyright Notice

Copyright (c) 2017 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
1.1. Terminology and Requirements Language	3
2. Bluetooth LE Networks and the IPSP	3
3. Specification of IPv6 mesh over Bluetooth LE networks	3
3.1. Protocol stack	4
3.2. Subnet model	4
3.3. Link model	5
3.3.1. Stateless address autoconfiguration	5
3.3.2. Neighbor Discovery	5
3.3.3. Header compression	6
3.3.4. Unicast and multicast mapping	7
4. IANA Considerations	8
5. Security Considerations	8
6. Acknowledgements	8
7. References	8
7.1. Normative References	9
7.2. Informative References	9
Authors' Addresses	10

1. Introduction

Bluetooth low energy (hereinafter, Bluetooth LE) was first introduced in the Bluetooth 4.0 specification. Bluetooth LE (which has been marketed as Bluetooth Smart) is a low-power wireless technology designed for short-range control and monitoring applications. Bluetooth LE is currently implemented in a wide range of consumer electronics devices, such as smartphones and wearable devices. Given the high potential of this technology for the Internet of Things, the Bluetooth Special Interest Group (Bluetooth SIG) and the IETF have produced specifications in order to enable IPv6 over Bluetooth LE, such as the Internet Protocol Support Profile (IPSP) [IPSP], and RFC 7668, respectively. Bluetooth 4.0 only supports Bluetooth LE networks that follow the star topology. In consequence, RFC 7668 was specifically developed and optimized for that type of network topology. However, subsequent Bluetooth specifications allow the formation of extended topologies [BTCorev4.1], such as the mesh topology. The functionality described in RFC 7668 is not sufficient and would fail to enable IPv6 over mesh networks composed of Bluetooth LE links. This document specifies the mechanisms needed to enable IPv6 over mesh networks composed of Bluetooth LE links. This specification also allows to run IPv6 over Bluetooth LE star topology

networks, albeit without all the topology-specific optimizations contained in RFC 7668.

1.1. Terminology and Requirements Language

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 RFC 2119 [RFC2119].

The terms 6LoWPAN Node (6LN), 6LoWPAN Router (6LR) and 6LoWPAN Border Router (6LBR) are defined as in [RFC6775], with an addition that Bluetooth LE central and Bluetooth LE peripheral (see Section 2) can both be adopted by a 6LN, a 6LR or a 6LBR.

2. Bluetooth LE Networks and the IPSP

Bluetooth LE defines two Generic Access Profile (GAP) roles of relevance herein: the Bluetooth LE central role and the Bluetooth LE peripheral role. A device in the central role, which is called central from now on, has traditionally been able to manage multiple simultaneous connections with a number of devices in the peripheral role, called peripherals hereinafter. Bluetooth 4.1 introduced the possibility for a peripheral to be connected to more than one central simultaneously, therefore allowing extended topologies beyond the star topology for a Bluetooth LE network. In addition, a device may simultaneously be a central in a set of link layer connections, as well as a peripheral in others. On the other hand, the IPSP enables discovery of IP-enabled devices and the establishment of a link layer connection for transporting IPv6 packets. The IPSP defines the Node and Router roles for devices that consume/originate IPv6 packets and for devices that can route IPv6 packets, respectively. Consistently with Bluetooth 4.1, a device may implement both roles simultaneously.

This document assumes a mesh network composed of Bluetooth LE links, where link layer connections have been established between neighboring IPv6-enabled devices. The IPv6 forwarding devices of the mesh have to implement both Node and Router roles, while simpler leaf-only nodes can implement only the Node role. In an IPv6-enabled mesh of Bluetooth LE links, a node is a neighbor of another node, and vice versa, if a link layer connection has been established between both by using the IPSP functionality for discovery and link layer connection establishment for IPv6 packet transport.

3. Specification of IPv6 mesh over Bluetooth LE networks

3.1. Protocol stack

Figure 1 illustrates the protocol stack for IPv6 mesh over Bluetooth LE networks. There are two main differences with the IPv6 over Bluetooth LE stack in RFC 7668: a) the adaptation layer below IPv6 (labelled as "6Lo for IPv6 mesh of Bluetooth LE") is now adapted for mesh networks of Bluetooth LE links, and b) the protocol stack for IPv6 mesh networks of Bluetooth LE links includes IPv6 routing functionality.

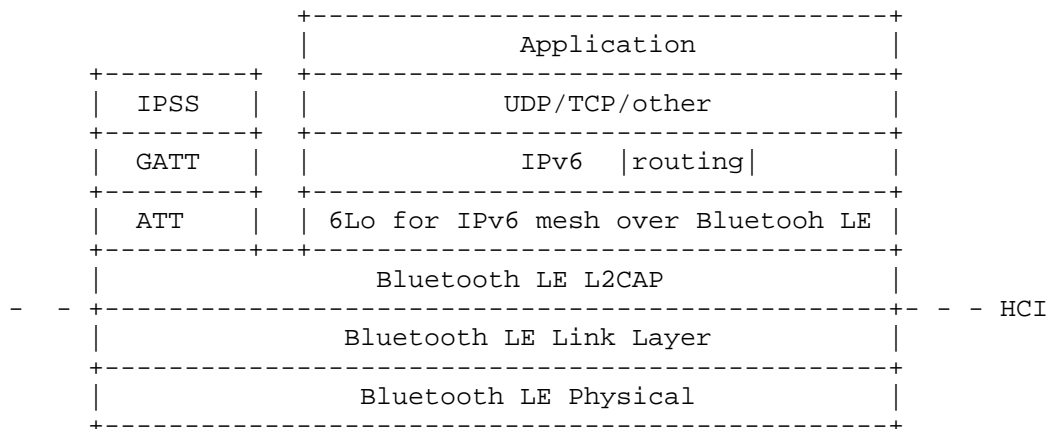


Figure 1: Protocol stack for IPv6 mesh over Bluetooth LE.

3.2. Subnet model

For IPv6 mesh over Bluetooth LE, a multilink model has been chosen, as further illustrated in Figure 2. As IPv6 over Bluetooth LE is intended for constrained nodes, and for Internet of Things use cases and environments, the complexity of implementing a separate subnet on each peripheral-central link and routing between the subnets appears to be excessive. In this specification, the benefits of treating the collection of point-to-point links between a central and its connected peripherals as a single multilink subnet rather than a multiplicity of separate subnets are considered to outweigh the multilink model's drawbacks as described in [RFC4903].




Figure 2: Example of an IPv6 mesh over a Bluetooth LE network connected to the Internet

One or more 6LBRs are connected to the Internet. 6LNs are connected to the network through a 6LR or a 6LBR. A prefix is used on the whole subnet.

IPv6 mesh networks over Bluetooth LE MUST follow a route-over approach. This document does not specify the routing protocol to be used in an IPv6 mesh over Bluetooth LE.

3.3. Link model

3.3.1. Stateless address autoconfiguration

6LN, 6LR and 6LBR IPv6 addresses in an IPv6 mesh over Bluetooth LE are configured as per section 3.2.2 of RFC 7668.

Multihop DAD functionality as defined in section 8.2 of RFC 6775, or some substitute mechanism (see section 3.3.2), MUST be supported.

3.3.2. Neighbor Discovery

'Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)' [RFC6775] describes the neighbor discovery approach as adapted for use in several 6LoWPAN topologies, including the mesh topology. The route-over functionality of RFC 6775 MUST be supported.

The following aspects of the Neighbor Discovery optimizations [RFC6775] are applicable to Bluetooth LE 6LNs:

1. A Bluetooth LE 6LN MUST NOT register its link-local address. A Bluetooth LE host MUST register its non-link-local addresses with its routers by sending a Neighbor Solicitation (NS) message with the Address Registration Option (ARO) and process the Neighbor Advertisement (NA) accordingly. The NS with the ARO option MUST be sent irrespective of the method used to generate the IID. The ARO option requires use of an EUI-64 identifier [RFC6775]. In the case of Bluetooth LE, the field SHALL be filled with the 48-bit device address used by the Bluetooth LE node converted into 64-bit Modified EUI-64 format [RFC4291].

If the 6LN registers for a same compression context multiple addresses that are not based on Bluetooth device address, the header compression efficiency will decrease.

2. For sending Router Solicitations and processing Router Advertisements the Bluetooth LE hosts MUST, respectively, follow Sections 5.3 and 5.4 of the [RFC6775].

3. The router behavior for 6LRs and 6LBRs is described in Section 6 of RFC 6775. However, as per this specification, routers SHALL NOT use multicast NSs to discover other routers' link layer addresses.

4. Border router behavior is described in Section 7 of RFC 6775.

RFC 6775 defines substitutable mechanisms for distributing prefixes and context information (section 8.1 of RFC 6775), as well as for Duplicate Address Detection across a route-over 6LoWPAN (section 8.2 of RFC 6775). Implementations of this specification MUST support the features described in sections 8.1 and 8.2 of RFC 6775 unless some alternative ("substitute") from some other specification is supported.

3.3.3. Header compression

Header compression as defined in RFC 6282 [RFC6282], which specifies the compression format for IPv6 datagrams on top of IEEE 802.15.4, is REQUIRED as the basis for IPv6 header compression on top of Bluetooth LE. All headers MUST be compressed according to RFC 6282 [RFC6282] encoding formats.

To enable efficient header compression, when the 6LBR sends a Router Advertisement it MUST include a 6LoWPAN Context Option (6CO) [RFC6775] matching each address prefix advertised via a Prefix Information Option (PIO) [RFC4861] for use in stateless address autoconfiguration.

The specific optimizations of RFC 7668 for header compression, which exploit the star topology and ARO, cannot be generalized in a mesh network composed of Bluetooth LE links. Still, a subset of those optimizations can be applied in some cases in such a network. In particular, the latter comprise link-local interactions, non-link-local packet transmissions originated and performed by a 6LN, and non-link-local packets transmitted (but not necessarily originated) by the neighbor of a 6LN to that 6LN. For the rest of packet transmissions, context-based compression MAY be used.

When a device transmits a packet to a neighbor, the sender MUST fully elide the source IID if the source IPv6 address is the link-local address based on the sender's Bluetooth device address (SAC=0, SAM=11). The sender also MUST fully elide the destination IPv6 address if it is the link-local-address based on the neighbor's Bluetooth device address (DAC=0, DAM=11).

When a 6LN transmits a packet, with a non-link-local source address that the 6LN has registered with ARO in the next-hop router for the indicated prefix, the source address MUST be fully elided if it is the latest address that the 6LN has registered for the indicated prefix (SAC=1, SAM=11). If the source non-link-local address is not the latest registered by the 6LN, then the 64-bits of the IID SHALL be fully carried in-line (SAC=1, SAM=01) or if the first 48-bits of the IID match with the latest address registered by the 6LN, then the last 16-bits of the IID SHALL be carried in-line (SAC=1, SAM=10).

When a router transmits a packet to a neighboring 6LN, with a non-link-local destination address, the router MUST fully elide the destination IPv6 address if the destination address is the latest registered by the 6LN with ARO for the indicated context (DAC=1, DAM=11). If the destination address is a non-link-local address and not the latest registered, then the 6LN MUST either include the IID part fully in-line (DAM=01) or, if the first 48-bits of the IID match to the latest registered address, then elide those 48-bits (DAM=10).

3.3.4. Unicast and multicast mapping

The Bluetooth LE Link Layer does not support multicast. Hence, traffic is always unicast between two Bluetooth LE neighboring nodes. If a node needs to send a multicast packet to several neighbors, it has to replicate the packet and unicast it on each link. However, this may not be energy efficient, and particular care must be taken if the node is battery powered. A router (i.e. a 6LR or a 6LBR) MUST keep track of neighboring multicast listeners, and it MUST NOT forward multicast packets to neighbors that have not registered as listeners for multicast groups the packets belong to.

4. IANA Considerations

There are no IANA considerations related to this document.

5. Security Considerations

The security considerations in RFC 7668 apply.

IPv6 mesh networks over Bluetooth LE require a routing protocol to find end-to-end paths. Unfortunately, the routing protocol may generate additional opportunities for threats and attacks to the network.

RFC 7416 [RFC 7416] provides a systematic overview of threats and attacks on the IPv6 Routing Protocol for Low-Power and Lossy Networks (RPL), as well as countermeasures. In that document, described threats and attacks comprise threats due to failures to authenticate, threats due to failure to keep routing information, threats and attacks on integrity, and threats and attacks on availability. Reported countermeasures comprise confidentiality attack, integrity attack, and availability attack countermeasures.

While this specification does not state the routing protocol to be used in IPv6 mesh over Bluetooth LE networks, the guidance of RFC 7416 is useful when RPL is used in such scenarios. Furthermore, such guidance may partly apply for other routing protocols as well.

6. Acknowledgements

The Bluetooth, Bluetooth Smart and Bluetooth Smart Ready marks are registered trademarks owned by Bluetooth SIG, Inc.

The authors of this document are grateful to all RFC 7668 authors, since this document borrows many concepts (albeit, with necessary extensions) from RFC 7668.

The authors also thank Alain Michaud, Mark Powell and Martin Turon for their comments, which helped improve the document.

Carles Gomez has been supported in part by the Spanish Government Ministerio de Economia y Competitividad through project TEC2012-32531, and FEDER.

7. References

7.1. Normative References

- [BTCorev4.1] Bluetooth Special Interest Group, "Bluetooth Core Specification Version 4.1", December 2013, <<https://www.bluetooth.org/en-us/specification/adopted-specifications>>.
- [IPSP] Bluetooth Special Interest Group, "Bluetooth Internet Protocol Support Profile Specification Version 1.0.0", December 2014, <<https://www.bluetooth.org/en-us/specification/adopted-specifications>>.
- [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>>.
- [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>>.
- [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, <<https://www.rfc-editor.org/info/rfc6282>>.
- [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>>.
- [RFC7668] Nieminen, J., Savolainen, T., Isomaki, M., Patil, B., Shelby, Z., and C. Gomez, "IPv6 over BLUETOOTH(R) Low Energy", RFC 7668, DOI 10.17487/RFC7668, October 2015, <<https://www.rfc-editor.org/info/rfc7668>>.

7.2. Informative References

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

[RFC7416] Tsao, T., Alexander, R., Dohler, M., Daza, V., Lozano, A., and M. Richardson, Ed., "A Security Threat Analysis for the Routing Protocol for Low-Power and Lossy Networks (RPLs)", RFC 7416, DOI 10.17487/RFC7416, January 2015, <<https://www.rfc-editor.org/info/rfc7416>>.

Authors' Addresses

Carles Gomez
Universitat Politecnica de Catalunya/Fundacio i2cat
C/Esteve Terradas, 7
Castelldefels 08860
Spain

Email: carlesgo@entel.upc.edu

Seyed Mahdi Darroudi
Universitat Politecnica de Catalunya/Fundacio i2cat
C/Esteve Terradas, 7
Castelldefels 08860
Spain

Email: sm.darroudi@entel.upc.edu

Teemu Savolainen
Nokia Technologies
Hatanpaan valtatie 30
Tampere 33100
Finland

Email: teemu.savolainen@nokia.com

6Lo Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 8, 2017

Y-H. Choi
Y-G. Hong
ETRI
J-S. Youn
Dongueui Univ
D-K. Kim
KNU
J-H. Choi
Samsung Electronics Co.,
March 7, 2017

Transmission of IPv6 Packets over Near Field Communication
draft-ietf-6lo-nfc-06

Abstract

Near field communication (NFC) is a set of standards for smartphones and portable devices to establish radio communication with each other by touching them together or bringing them into proximity, usually no more than 10 cm. NFC standards cover communications protocols and data exchange formats, and are based on existing radio-frequency identification (RFID) standards including ISO/IEC 14443 and FeliCa. The standards include ISO/IEC 18092 and those defined by the NFC Forum. The NFC technology has been widely implemented and available in mobile phones, laptop computers, and many other devices. This document describes how IPv6 is transmitted over NFC using 6LowPAN techniques.

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 September 8, 2017.

Copyright Notice

Copyright (c) 2017 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	3
2. Conventions and Terminology	3
3. Overview of Near Field Communication Technology	4
3.1. Peer-to-peer Mode of NFC	4
3.2. Protocol Stacks of NFC	4
3.3. NFC-enabled Device Addressing	6
3.4. NFC MAC PDU Size and MTU	6
4. Specification of IPv6 over NFC	7
4.1. Protocol Stacks	7
4.2. Link Model	7
4.3. Stateless Address Autoconfiguration	8
4.4. IPv6 Link Local Address	9
4.5. Neighbor Discovery	9
4.6. Dispatch Header	10
4.7. Header Compression	10
4.8. Fragmentation and Reassembly	11
4.9. Unicast Address Mapping	11
4.10. Multicast Address Mapping	12
5. Internet Connectivity Scenarios	12
5.1. NFC-enabled Device Connected to the Internet	12
5.2. Isolated NFC-enabled Device Network	13
6. IANA Considerations	13
7. Security Considerations	13
8. Acknowledgements	14
9. References	14
9.1. Normative References	14
9.2. Informative References	15
Authors' Addresses	16

1. Introduction

NFC is a set of short-range wireless technologies, typically requiring a distance of 10 cm or less. NFC operates at 13.56 MHz on ISO/IEC 18000-3 air interface and at rates ranging from 106 kbit/s to 424 kbit/s. NFC always involves an initiator and a target; the initiator actively generates an RF field that can power a passive target. This enables NFC targets to take very simple form factors such as tags, stickers, key fobs, or cards that do not require batteries. NFC peer-to-peer communication is possible, provided both devices are powered. NFC builds upon RFID systems by allowing two-way communication between endpoints, where earlier systems such as contactless smart cards were one-way only. It has been used in devices such as mobile phones, running Android operating system, named with a feature called "Android Beam". In addition, it is expected for the other mobile phones, running the other operating systems (e.g., iOS, etc.) to be equipped with NFC technology in the near future.

Considering the potential for exponential growth in the number of heterogeneous air interface technologies, NFC would be widely used as one of the other air interface technologies, such as Bluetooth Low Energy (BT-LE), Wi-Fi, and so on. Each of the heterogeneous air interface technologies has its own characteristics, which cannot be covered by the other technologies, so various kinds of air interface technologies would co-exist together. Therefore, it is required for them to communicate with each other. NFC also has the strongest ability (e.g., secure communication distance of 10 cm) to prevent a third party from attacking privacy.

When the number of devices and things having different air interface technologies communicate with each other, IPv6 is an ideal internet protocols owing to its large address space. Also, NFC would be one of the endpoints using IPv6. Therefore, this document describes how IPv6 is transmitted over NFC using 6LoWPAN techniques.

RFC4944 [1] specifies the transmission of IPv6 over IEEE 802.15.4. The NFC link also has similar characteristics to that of IEEE 802.15.4. Many of the mechanisms defined in RFC 4944 [1] can be applied to the transmission of IPv6 on NFC links. This document specifies the details of IPv6 transmission over NFC links.

2. Conventions and 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 [2].

3. Overview of Near Field Communication Technology

NFC technology enables simple and safe two-way interactions between electronic devices, allowing consumers to perform contactless transactions, access digital content, and connect electronic devices with a single touch. NFC complements many popular consumer level wireless technologies, by utilizing the key elements in existing standards for contactless card technology (ISO/IEC 14443 A&B and JIS-X 6319-4). NFC can be compatible with existing contactless card infrastructure and it enables a consumer to utilize one device across different systems.

Extending the capability of contactless card technology, NFC also enables devices to share information at a distance that is less than 10 cm with a maximum communication speed of 424 kbps. Users can share business cards, make transactions, access information from a smart poster or provide credentials for access control systems with a simple touch.

NFC's bidirectional communication ability is ideal for establishing connections with other technologies by the simplicity of touch. In addition to the easy connection and quick transactions, simple data sharing is also available.

3.1. Peer-to-peer Mode of NFC

NFC-enabled devices are unique in that they can support three modes of operation: card emulation, peer-to-peer, and reader/writer. Peer-to-peer mode enables two NFC-enabled devices to communicate with each other to exchange information and share files, so that users of NFC-enabled devices can quickly share contact information and other files with a touch. Therefore, an NFC-enabled device can securely send IPv6 packets to any corresponding node on the Internet when an NFC-enabled gateway is linked to the Internet.

3.2. Protocol Stacks of NFC

IP can use the services provided by the Logical Link Control Protocol (LLCP) in the NFC stack to provide reliable, two-way transport of information between the peer devices. Figure 1 depicts the NFC P2P protocol stack with IPv6 bindings to LLCP.

For data communication in IPv6 over NFC, an IPv6 packet SHALL be passed down to LLCP of NFC and transported to an Information Field in Protocol Data Unit (I PDU) of LLCP of the NFC-enabled peer device. LLCP does not support fragmentation and reassembly. For IPv6 addressing or address configuration, LLCP SHALL provide related information, such as link layer addresses, to its upper layer. The

LLCP to IPv6 protocol binding SHALL transfer the SSAP and DSAP value to the IPv6 over NFC protocol. SSAP stands for Source Service Access Point, which is a 6-bit value meaning a kind of Logical Link Control (LLC) address, while DSAP means an LLC address of the destination NFC-enabled device.




Figure 1: Protocol Stacks of NFC

The LLCP consists of Logical Link Control (LLC) and MAC Mapping. The MAC Mapping integrates an existing RF protocol into the LLCP architecture. The LLC contains three components, such as Link Management, Connection-oriented Transport, and Connection-less Transport. The Link Management component is responsible for serializing all connection-oriented and connection-less LLC PDU (Protocol Data Unit) exchanges and for aggregation and disaggregation of small PDUs. This component also guarantees asynchronous balanced mode communication and provides link status supervision by performing the symmetry procedure. The Connection-oriented Transport component is responsible for maintaining all connection-oriented data exchanges including connection set-up and termination. The Connectionless Transport component is responsible for handling unacknowledged data exchanges.

3.3. NFC-enabled Device Addressing

According to NFCForum-TS-LLCP_1.3 [3], NFC-enabled devices have two types of 6-bit addresses (i.e., SSAP and DSAP) to identify service access points. The several service access points can be installed on a NFC device. However, the SSAP and DSAP can be used as identifiers for NFC link connections with the IPv6 over NFC adaptation layer. Therefore, the SSAP can be used to generate an IPv6 interface identifier. Address values between 00h and 0Fh of SSAP and DSAP are reserved for identifying the well-known service access points, which are defined in the NFC Forum Assigned Numbers Register. Address values between 10h and 1Fh SHALL be assigned by the local LLC to services registered by local service environment. In addition, address values between 20h and 3Fh SHALL be assigned by the local LLC as a result of an upper layer service request. Therefore, the address values between 20h and 3Fh can be used for generating IPv6 interface identifiers.

3.4. NFC MAC PDU Size and MTU

As mentioned in Section 3.2, an IPv6 packet SHALL be passed down to LLC of NFC and transported to an Unnumbered Information Protocol Data Unit (UI PDU) and an Information Field in Protocol Data Unit (I PDU) of LLC of the NFC-enabled peer device.

The information field of an I PDU SHALL contain a single service data unit. The maximum number of octets in the information field is determined by the Maximum Information Unit (MIU) for the data link connection. The default value of the MIU for I PDUs SHALL be 128 octets. The local and remote LLCs each establish and maintain distinct MIU values for each data link connection endpoint. Also, an LLC MAY announce a larger MIU for a data link connection by transmitting an MIUX extension parameter within the information field. If no MIUX parameter is transmitted, the default MIU value of 128 SHALL be used. Otherwise, the MTU size in NFC LLC SHALL calculate the MIU value as follows:

$$\text{MIU} = 128 + \text{MIUX}.$$

When the MIUX parameter is encoded as a TLV, the TLV Type field SHALL be 0x02 and the TLV Length field SHALL be 0x02. The MIUX parameter SHALL be encoded into the least significant 11 bits of the TLV Value field. The unused bits in the TLV Value field SHALL be set to zero by the sender and SHALL be ignored by the receiver. However, a maximum value of the TLV Value field can be 0x7FF, and a maximum size of the MTU in NFC LLC is 2176 bytes.

4. Specification of IPv6 over NFC

NFC technology also has considerations and requirements owing to low power consumption and allowed protocol overhead. 6LoWPAN standards RFC 4944 [1], RFC 6775 [4], and RFC 6282 [5] provide useful functionality for reducing overhead which can be applied to NFC. This functionality consists of link-local IPv6 addresses and stateless IPv6 address auto-configuration (see Section 4.3), Neighbor Discovery (see Section 4.5) and header compression (see Section 4.7).

4.1. Protocol Stacks

Figure 2 illustrates IPv6 over NFC. Upper layer protocols can be transport layer protocols (TCP and UDP), application layer protocols, and others capable running on top of IPv6.

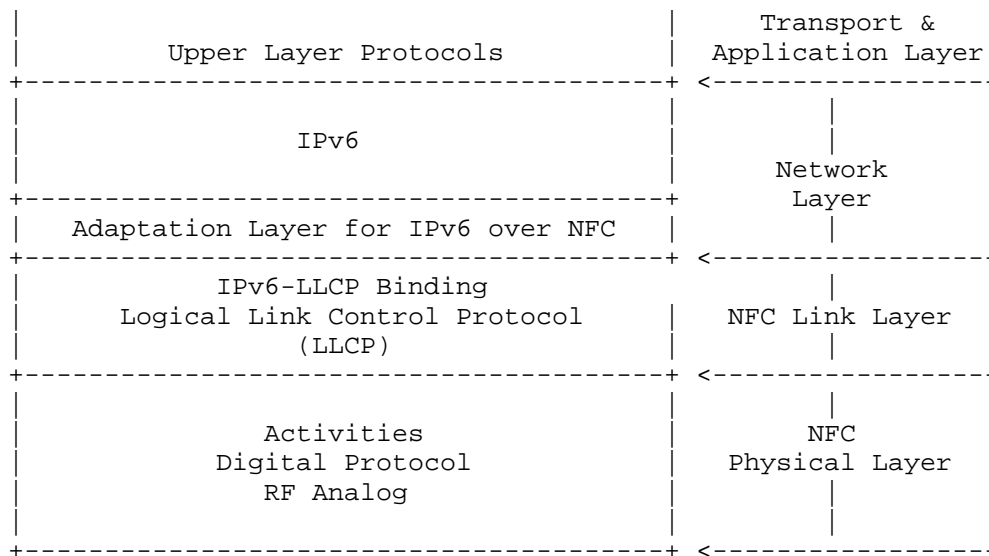


Figure 2: Protocol Stacks for IPv6 over NFC

The adaptation layer for IPv6 over NFC SHALL support neighbor discovery, stateless address auto-configuration, header compression, and fragmentation & reassembly.

4.2. Link Model

In the case of BT-LE, the Logical Link Control and Adaptation Protocol (L2CAP) supports fragmentation and reassembly (FAR) functionality; therefore, the adaptation layer for IPv6 over BT-LE does not have to conduct the FAR procedure. The NFC LLCP, in

contrast, does not support the FAR functionality, so IPv6 over NFC needs to consider the FAR functionality, defined in RFC 4944 [1]. However, the MTU on an NFC link can be configured in a connection procedure and extended enough to fit the MTU of IPv6 packet (see Section 4.8).

The NFC link between two communicating devices is considered to be a point-to-point link only. Unlike in BT-LE, an NFC link does not support a star topology or mesh network topology but only direct connections between two devices. Furthermore, the NFC link layer does not support packet forwarding in link layer. Due to this characteristics, 6LoWPAN functionalities, such as addressing and auto-configuration, and header compression, need to be specialized into IPv6 over NFC.

4.3. Stateless Address Autoconfiguration

An NFC-enabled device (i.e., 6LN) performs stateless address autoconfiguration as per RFC 4862 [6]. A 64-bit Interface identifier (IID) for an NFC interface is formed by utilizing the 6-bit NFC LLCP address (see Section 3.3). In the viewpoint of address configuration, such an IID SHOULD guarantee a stable IPv6 address because each data link connection is uniquely identified by the pair of DSAP and SSAP included in the header of each LLC PDU in NFC.

Following the guidance of RFC 7136 [10], interface identifiers of all unicast addresses for NFC-enabled devices are 64 bits long and constructed in a modified EUI-64 format as shown in Figure 3.




Figure 3: Formation of IID from NFC-enabled device address

The 'R' bits are output values which MAY be created by mechanisms like hash functions with input values, i.e., the SSAP and other values (e.g., prefix) because the 6-bit address of SSAP is easy and short to be targeted by attacks of third party (e.g., address scanning). Figure 4 shows an example for IID creation. The F() means a mechanism to make a output value for 64-bit IID, and an parameter, "offset" is an example input value for making the different output values.

$IID = F(\text{SHA-256}(6\text{-bit SSAP}, 64\text{-bit Prefix}), 'u' \text{ bit}, \text{offset})$

Figure 4: An example of an IID creation mechanism

In addition, the "Universal/Local" bit (i.e., the 'u' bit) of an NFC-enabled device address MUST be set to 0 RFC 4291 [7].

4.4. IPv6 Link Local Address

Only if the NFC-enabled device address is known to be a public address, the "Universal/Local" bit be set to 1. The IPv6 link-local address for an NFC-enabled device is formed by appending the IID, to the prefix FE80::/64, as depicted in Figure 5.




Figure 5: IPv6 link-local address in NFC

The tool for a 6LBR to obtain an IPv6 prefix for numbering the NFC network is can be accomplished via DHCPv6 Prefix Delegation (RFC 3633 [8]).

4.5. Neighbor Discovery

Neighbor Discovery Optimization for 6LoWPANs (RFC 6775 [4]) describes the neighbor discovery approach in several 6LoWPAN topologies, such as mesh topology. NFC does not support a complicated mesh topology but only a simple multi-hop network topology or directly connected peer-to-peer network. Therefore, the following aspects of RFC 6775 are applicable to NFC:

1. In a case that an NFC-enabled device (6LN) is directly connected to a 6LBR, an NFC 6LN MUST register its address with the 6LBR by sending a Neighbor Solicitation (NS) message with the Address Registration Option (ARO) and process the Neighbor Advertisement (NA) accordingly. In addition, if DHCPv6 is used to assign an address, Duplicate Address Detection (DAD) MAY not be required.

2. For sending Router Solicitations and processing Router Advertisements the NFC 6LNs MUST follow Sections 5.3 and 5.4 of RFC 6775.

4.6. Dispatch Header

All IPv6-over-NFC encapsulated datagrams are prefixed by an encapsulation header stack consisting of a Dispatch value followed by zero or more header fields. The only sequence currently defined for IPv6-over-NFC is the LOWPAN_IPHC header followed by payload, as depicted in Figure 6.




Figure 6: A IPv6-over-NFC Encapsulated 6LOWPAN_IPHC Compressed IPv6 Datagram

The dispatch value may be treated as an unstructured namespace. Only a single pattern is used to represent current IPv6-over-NFC functionality.

+-----+	+-----+	+-----+
Pattern	Header Type	Reference
+-----+	+-----+	+-----+
01 1xxxxx	6LOWPAN_IPHC	[RFC6282]
+-----+	+-----+	+-----+

Figure 7: Dispatch Values

Other IANA-assigned 6LoWPAN Dispatch values do not apply to this specification.

4.7. Header Compression

Header compression as defined in RFC 6282 [5], which specifies the compression format for IPv6 datagrams on top of IEEE 802.15.4, is REQUIRED in this document as the basis for IPv6 header compression on top of NFC. All headers MUST be compressed according to RFC 6282 encoding formats.

Therefore, IPv6 header compression in RFC 6282 [5] MUST be implemented. Further, implementations MAY also support Generic Header Compression (GHC) of RFC 7400 [11].

If a 16-bit address is required as a short address, it MUST be formed by padding the 6-bit NFC link-layer (node) address to the left with zeros as shown in Figure 8.




Figure 8: NFC short address format

4.8. Fragmentation and Reassembly

NFC provides fragmentation and reassembly (FAR) for payloads from 128 bytes up to 2176 bytes as mentioned in Section 3.4. The MTU of a general IPv6 packet can fit into a single NFC link frame. Therefore, the FAR functionality as defined in RFC 4944, which specifies the fragmentation methods for IPv6 datagrams on top of IEEE 802.15.4, MAY NOT be required as the basis for IPv6 datagram FAR on top of NFC. The NFC link connection for IPv6 over NFC MUST be configured with an equivalent MIU size to fit the MTU of IPv6 Packet. If NFC devices support extension of the MTU, the MIUX value is 0x480 in order to fit the MTU (1280 bytes) of a IPv6 packet.

4.9. Unicast Address Mapping

The address resolution procedure for mapping IPv6 non-multicast addresses into NFC link-layer addresses follows the general description in Section 7.2 of RFC 4861 [9], unless otherwise specified.

The Source/Target link-layer Address option has the following form when the addresses are 6-bit NFC link-layer (node) addresses.




Figure 9: Unicast address mapping

Option fields:

Type:

1: for Source Link-layer address.

2: for Target Link-layer address.

Length:

This is the length of this option (including the type and length fields) in units of 8 octets. The value of this field is 1 for 6-bit NFC node addresses.

NFC address:

The 6-bit address in canonical bit order. This is the unicast address the interface currently responds to.

4.10. Multicast Address Mapping

All IPv6 multicast packets MUST be sent to NFC Destination Address, 0x3F (broadcast) and be filtered at the IPv6 layer. When represented as a 16-bit address in a compressed header, it MUST be formed by padding on the left with a zero. In addition, the NFC Destination Address, 0x3F, MUST NOT be used as a unicast NFC address of SSAP or DSAP.

```

      0                               1
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-----+-----+-----+-----+
| Padding(all zeros) | 1 1 1 1 1 1 |
+-----+-----+-----+-----+

```

Figure 10: Multicast address mapping

5. Internet Connectivity Scenarios

As two typical scenarios, the NFC network can be isolated and connected to the Internet.

5.1. NFC-enabled Device Connected to the Internet

One of the key applications of using IPv6 over NFC is securely transmitting IPv6 packets because the RF distance between 6LN and 6LBR is typically within 10 cm. If any third party wants to hack into the RF between them, it must come to nearly touch them. Applications can choose which kinds of air interfaces (e.g., BT-LE,

Wi-Fi, NFC, etc.) to send data depending on the characteristics of the data.

Figure 11 illustrates an example of an NFC-enabled device network connected to the Internet. The distance between 6LN and 6LBR is typically 10 cm or less. If there is any laptop computers close to a user, it will become the a 6LBR. Additionally, when the user mounts an NFC-enabled air interface adapter (e.g., portable NFC dongle) on the close laptop PC, the user's NFC-enabled device (6LN) can communicate with the laptop PC (6LBR) within 10 cm distance.



Figure 11: NFC-enabled device network connected to the Internet

5.2. Isolated NFC-enabled Device Network

In some scenarios, the NFC-enabled device network may transiently be a simple isolated network as shown in the Figure 12.




Figure 12: Isolated NFC-enabled device network

In mobile phone markets, applications are designed and made by user developers. They may image interesting applications, where three or more mobile phones touch or attach each other to accomplish outstanding performance.

6. IANA Considerations

There are no IANA considerations related to this document.

7. Security Considerations

When interface identifiers (IIDs) are generated, devices and users are required to consider mitigating various threats, such as correlation of activities over time, location tracking, device-specific vulnerability exploitation, and address scanning.

IPv6-over-NFC is, in practice, not used for long-lived links for big size data transfer or multimedia streaming, but used for extremely short-lived links (i.e., single touch-based approaches) for ID verification and mobile payment. This will mitigate the threat of correlation of activities over time.

IPv6-over-NFC uses an IPv6 interface identifier formed from a "Short Address" and a set of well-known constant bits (such as padding with '0's) for the modified EUI-64 format. However, the short address of NFC link layer (LLC) is not generated as a physically permanent value but logically generated for each connection. Thus, every single touch connection can use a different short address of NFC link with an extremely short-lived link. This can mitigate address scanning as well as location tracking and device-specific vulnerability exploitation.

However, malicious tries for one connection of a long-lived link with NFC technology are not secure, so the method of deriving interface identifiers from 6-bit NFC Link layer addresses is intended to preserve global uniqueness when it is possible. Therefore, it requires a way to protect from duplication through accident or forgery and to define a way to include sufficient bit of entropy in the IPv6 interface identifier, such as random EUI-64.

8. Acknowledgements

We are grateful to the members of the IETF 6lo working group.

Michael Richardson, Suresh Krishnan, Pascal Thubert, Carsten Bormann, and Alexandru Petrescu have provided valuable feedback for this draft.

9. References

9.1. Normative References

- [1] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<http://www.rfc-editor.org/info/rfc4944>>.
- [2] 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>>.
- [3] "NFC Logical Link Control Protocol version 1.3", NFC Forum Technical Specification , March 2016.

- [4] 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>>.
- [5] 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>>.
- [6] 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>>.
- [7] 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>>.
- [8] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, DOI 10.17487/RFC3633, December 2003, <<http://www.rfc-editor.org/info/rfc3633>>.
- [9] 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>>.
- [10] Carpenter, B. and S. Jiang, "Significance of IPv6 Interface Identifiers", RFC 7136, DOI 10.17487/RFC7136, February 2014, <<http://www.rfc-editor.org/info/rfc7136>>.
- [11] Bormann, C., "6LoWPAN-GHC: Generic Header Compression for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 7400, DOI 10.17487/RFC7400, November 2014, <<http://www.rfc-editor.org/info/rfc7400>>.

9.2. Informative References

- [12] "Near Field Communication - Interface and Protocol (NFCIP-1) 3rd Ed.", ECMA-340 , June 2013.

Authors' Addresses

Younghwan Choi
Electronics and Telecommunications Research Institute
218 Gajeongno, Yuseong
Daejeon 305-700
Korea

Phone: +82 42 860 1429
Email: yhc@etri.re.kr

Yong-Geun Hong
Electronics and Telecommunications Research Institute
161 Gajeong-Dong Yuseung-Gu
Daejeon 305-700
Korea

Phone: +82 42 860 6557
Email: yghong@etri.re.kr

Joo-Sang Youn
DONG-EUI University
176 Eomgwangno Busan_jin_gu
Busan 614-714
Korea

Phone: +82 51 890 1993
Email: joosang.youn@gmail.com

Dongkyun Kim
Kyungpook National University
80 Daehak-ro, Buk-gu
Daegu 702-701
Korea

Phone: +82 53 950 7571
Email: dongkyun@knu.ac.kr

JinHyouk Choi
Samsung Electronics Co.,
129 Samsung-ro, Youngdong-gu
Suwon 447-712
Korea

Phone: +82 2 2254 0114
Email: jinchoe@samsung.com

6lo
Internet-Draft
Updates: 6775 (if approved)
Intended status: Standards Track
Expires: July 14, 2017

P. Thubert, Ed.
cisco
E. Nordmark
Arista Networks
S. Chakrabarti
January 10, 2017

An Update to 6LoWPAN ND
draft-ietf-6lo-rfc6775-update-01

Abstract

This specification updates 6LoWPAN Neighbor Discovery (RFC6775), to clarify the role of the protocol as a registration technique, simplify the registration operation in 6LoWPAN routers, and provide enhancements to the registration capabilities, in particular for the registration to a backbone router for proxy ND operations.

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 July 14, 2017.

Copyright Notice

Copyright (c) 2017 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. Terminology	3
3. Updating RFC 6775	4
3.1. Transaction ID	4
3.2. Owner Unique ID	5
3.3. Extended Address Registration Option	5
3.4. Registering the Target Address	6
3.5. Link-local Addresses and Registration	6
4. Applicability and Requirements Served	8
5. The Enhanced Address Registration Option (EARO)	8
6. Backward Compatibility	12
6.1. Legacy 6LoWPAN Node	12
6.2. Legacy 6LoWPAN Router	12
6.3. Legacy 6LoWPAN Border Router	13
7. Security Considerations	13
8. IANA Considerations	14
9. Acknowledgments	14
10. References	14
10.1. Normative References	14
10.2. Informative References	15
10.3. External Informative References	17
Appendix A. Requirements	18
A.1. Requirements Related to Mobility	18
A.2. Requirements Related to Routing Protocols	18
A.3. Requirements Related to the Variety of Low-Power Link types	19
A.4. Requirements Related to Proxy Operations	20
A.5. Requirements Related to Security	20
A.6. Requirements Related to Scalability	22
Authors' Addresses	22

1. Introduction

IPv6 Neighbor Discovery (ND) Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs) [RFC6775] introduced a proactive registration mechanism to IPv6 ND services that is well suited to nodes belonging to a LLN.

The scope of this draft is an IPv6 Low Power Lossy Network (LLN), which can be a simple star or a more complex mesh topology. The LLN may be anchored at an IPv6 Backbone Router (6BBR). The Backbone Routers interconnect the LLNs over a Backbone Link and emulate that the LLN nodes are present on the Backbone using proxy-ND operations.

This specification modifies and extends the behaviour and protocol elements of [RFC6775] to enable additional capabilities, in particular the registration to a 6BBR for proxy ND operations [I-D.ietf-6lo-backbone-router].

2. 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].

Additionally, this document uses terminology from "Terms Used in Routing for Low-Power and Lossy Networks" [RFC7102] and [I-D.ietf-6tisch-terminology], as well as this additional terminology:

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 6BBR 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.

3. Updating RFC 6775

The support of this specification is signaled in Router Advertisement (RA) messages by 6LoWPAN Router (6LR) (how: tbd). Support for this specification can also be inferred from the update of the ARO option in the ND exchanges.

A Registering Node that supports this specification will favor registering to a 6LR that indicates support for this specification over that of [RFC6775].

3.1. 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.

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.

3.2. 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.ietf-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.

3.3. Extended Address Registration Option

This specification extends the Address Registration Option (ARO) used for the process of address registration. The new ARO is referred to as Extended ARO (EARO), and its semantics are modified as follows:

The address that is being registered with a Neighbor Solicitation (NS) with an EARO is now the Target Address, as opposed to the Source Address as specified in [RFC6775]. This change enables a 6LBR to use an address of his as source to the proxy-registration of an address that belongs to a LLN Node to a 6BBR. This also limits the use of an address as source address before it is registered and the associated Duplicate Address Detection (DAD) is complete.

The Unique ID in the EARO option does no more have to be a MAC address. A new TLV format is introduced and a IANA registry is created for the type (TBD). This enables in particular the use of a Provable Temporary UID (PT-UID) as opposed to burn-in MAC address, the PT-UID providing a trusted anchor by the 6LR and 6LBR to protect the state associated to the node.

The specification introduces a Transaction ID (TID) field in the EARO. The TID MUST be provided by a node that supports this specification and a new T flag MUST be set to indicate so. The T bit can be used to determine whether the peer supports this specification.

3.4. Registering the Target Address

One of the requirements that this specification serves is the capability by a router such as a RPL root to proxy-register an address to a 6BBR on behalf of a 6LN, as discussed in Appendix A.4. In order to serve that requirement, this specification changes the behaviour of the 6LN and the 6LR so that the Registered Address is found in the Target Address field of the NS and NA messages as opposed to the Source Address.

With this convention, a TLLA option would indicate the link-layer address of the 6LN that owns the address, whereas the SLLA Option in a NS message indicates that of the Registering Node, which can be the owner device, or a proxy.

Since the Registering Node is the one that has reachability with the 6LR, and is the one expecting packets for the 6LN, it makes sense to maintain compatibility with [RFC6775], and it is REQUIRED that an SLLA Option is always placed in a registration NS(EARO) message.

3.5. Link-local Addresses and Registration

Considering that LLN nodes are often not wired and may move, there is no guarantee that a link-local address stays unique between a potentially variable and unbounded set of neighboring nodes. Compared to [RFC6775], this specification only requires that a link-local address is unique from the perspective of the peering nodes. This simplifies the Duplicate Address Detection (DAD) for link-local addresses, and there is no DAR/DAC exchange between the 6LR and a 6LBR for link-local addresses.

Additionally, [RFC6775] requires that a 6LoWPAN Node (6LN) uses an address being registered as the source of the registration message. This generates complexities in the 6LR to be able to cope with a potential duplication, in particular for global addresses. To simplify this, a 6LN and a 6LR that conform this specification always use link-local addresses as source and destination addresses for the registration NS/NA exchange. As a result, the registration is globally faster, and some of the complexity is removed.

In more details:

An exchange between two nodes using link-local addresses implies that they are reachable over one hop and that at least one of the 2 nodes acts as a 6LR. A node MUST register a link-local address to a 6LR in order to obtain reachability from that 6LR beyond the current exchange, and in particular to use the link-local address as source address to register other addresses, e.g. global addresses. If there is no collision with an address previously registered to this 6LR by another 6LN, then, from the standpoint of this 6LR, this link-local address is unique and the registration is acceptable. Conversely, it may possibly happen that two different 6LRs expose a same link-local address but different link-layer addresses. In that case, a 6LN may only interact with one of the 6LR so as to avoid confusion in the 6LN neighbor cache.

The DAD process between the 6LR and a 6LoWPAN Border Router (6LBR), which is based on a Duplicate Address Request (DAR) / Duplicate Address Confirmation (DAC) exchange as described in [RFC6775], does not need to take place for link-local addresses.

It is desired that a 6LR does not need to modify its state associated to the Source Address of an NS(EARO) message. For that reason, when possible, it is RECOMMENDED to use an address that is already registered with a 6LR

When registering to a 6LR that conforms this specification, a node MUST use a link-local address as the source address of the registration, whatever the type of IPv6 address that is being registered. That link-local Address MUST be either already registered, or the address that is being registered.

When a Registering Node does not have an already-registered address, it MUST register a link-local address, using it as both the Source and the Target Address of an NS(EARO) message. In that case, it is RECOMMENDED to use a link-local address that is (expected to be) globally unique, e.g. derived from a burn-in MAC address. An EARO option in the response NA indicates that the 6LR supports this specification.

Since there is no DAR/DAC exchange for link-local addresses, the 6LR may answer immediately to the registration of a link-local address, based solely on its existing state and the Source Link-Layer Option that MUST be placed in the NS(EARO) message as required in [RFC6775].

A node needs to register its IPv6 Global Unicast IPv6 Addresses (GUA) to a 6LR in order to obtain a global reachability for these addresses via that 6LR. As opposed to a node that complies to [RFC6775], a Registering Node registering a GUA does not use that GUA as Source Address for the registration to a 6LR that conforms this

specification. The DAR/DAC exchange MUST take place for non-link-local addresses as prescribed by [RFC6775].

4. Applicability and Requirements Served

This specification extends 6LoWPAN ND to sequence the registration and serves the requirements expressed Appendix A.1 by enabling the mobility of devices from one LLN to the next based on the complementary work in [I-D.ietf-6lo-backbone-router].

In the context of the the TimeSlotted Channel Hopping (TSCH) mode of [IEEEstd802154], 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.

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 discussed in Appendix A.3

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.

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 IEEE std 802.15.4 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 serves scalability requirements listed in Appendix A.6.

5. 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 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 1: EARO

Option Fields

Type:

Length: 2

Status:

Value	Description
0..2	See [RFC6775]. Note that a Status of 1 "Duplicate Address" applies to the Registered Address. If the Source Address conflicts with an existing registration, "Duplicate Source Address" should be used instead
3	Moved: The registration fails because it is not the freshest
4	Removed: The binding state was removed. This may be placed in an asynchronous NS(ARO) message, or as the rejection of a proxy registration to a Backbone Router
5	Proof requested: The registering node is challenged for owning the registered address or for being an acceptable proxy for the registration
6	Duplicate Source Address: The address used as source of the NS(ARO) conflicts with an existing registration.
7	Administrative Rejection: The address being registered is reserved for another use by an administrative decision (e.g. placed in a DHCPv6 pool); The Registering Node is requested to form a different address and retry
8	Invalid Registered Address: The address being registered is not usable on this link, e.g. it is not topologically correct
9	Invalid Source Address: The address used as source of the NS(ARO) is not usable on this link, e.g. it is not topologically correct

Table 1

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: 16-bit integer; expressed in minutes. 0 means that the registration has ended and the state should be removed.

Owner Unique Identifier (OUI): 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.

New status values are introduced, their values to be confirmed by IANA:

Moved: This status indicates that the registration is rejected because another more recent registration was done, as indicated by a same OUI and a more recent TID. One possible cause is a stale registration that has progressed slowly in the network and was passed by a more recent one. It could also indicate a OUI collision.

Removed: This status is expected in asynchronous messages from a registrar (6LR, 6LBR, 6BBR) to indicate that the registration state is removed, for instance due to time out of a lifetime, or a movement. It is used for instance by a 6BBR in a NA(ARO) message to indicate that the ownership of the proxy state on the backbone was transferred to another 6BBR, which is indicative of a movement of the device. The receiver of the NA is the device that has performed a registration that is now stale and it should clean up its state.

6. Backward Compatibility

6.1. Legacy 6LoWPAN Node

A legacy 6LN will use the registered address as source and will not use an EARO option. In order to be backward compatible, an updated 6LR needs to accept that registration if it is valid per [RFC3972], and manage the binding cache accordingly.

The main difference with [RFC3972] is that DAR/DAC exchange for DAD may be avoided for link-local addresses. Additionally, the 6LR SHOULD use an EARO in the reply, and may use all the status codes defined in this specification.

6.2. Legacy 6LoWPAN Router

The first registration by a an updated 6LN is for a link-local address, using that link-local address as source. A legacy 6LN will not makes a difference and accept -or reject- that registration as if the 6LN was a legacy node.

An updated 6LN will always use an EARO option in the registration NS message, whereas a legacy 6LN will always reply with an ARO option in the NA message. So from that first registration, the updated 6LN can figure whether the 6LR supports this specification or not.

When facing a legacy 6LR, an updated 6LN may attempt to find an alternate 6LR that is updated. In order to be backward compatible, based on the discovery that a 6LR is legacy, the 6LN needs to fallback to legacy behaviour and source the packet with the registered address.

The main difference is that the updated 6LN SHOULD use an EARO in the request regardless of the type of 6LN, legacy or updated

6.3. Legacy 6LoWPAN Border Router

With this specification, the DAR/DAC transports an EARO option as opposed to an ARO option. As described for the NS/NA exchange, devices that support this specification always use an EARO option and all the associated behaviour.

7. 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.ietf-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.

The LLN nodes depend on the 6LBR and the 6BBR for their operation. A trust model must be put in place to ensure that the right devices are acting in these roles, so as to avoid threats such as black-holing, or bombing attack whereby an impersonated 6LBR would destroy state in the network by using the "Removed" status code.

8. IANA Considerations

This document requires the following additions:

Address Registration Option Status Values Registry

Status	Description
3	Moved
4	Removed
5	Proof requested
6	Invalid Source Address
7	Administrative Rejection

IANA is required to change the registry accordingly

Table 2: New ARO Status values

9. Acknowledgments

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

10. References

10.1. 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>>.
- [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>>.
- [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>>.

10.2. 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-06 (work in progress), October 2016.
- [I-D.ietf-6lo-ap-nd]
Sarikaya, B., Thubert, P., and M. Sethi, "Address Protected Neighbor Discovery for Low-power and Lossy Networks", draft-ietf-6lo-ap-nd-00 (work in progress), November 2016.
- [I-D.ietf-6lo-backbone-router]
Thubert, P., "IPv6 Backbone Router", draft-ietf-6lo-backbone-router-02 (work in progress), September 2016.

- [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-09 (work in progress), December 2016.
- [I-D.ietf-6lo-nfc]
Choi, Y., Youn, J., and Y. Hong, "Transmission of IPv6 Packets over Near Field Communication", draft-ietf-6lo-nfc-05 (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-10 (work in progress), June 2016.
- [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-08 (work in progress), December 2016.
- [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-05 (work in progress), October 2016.
- [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.popa-6lo-6loplc-ipv6-over-ieee19012-networks]
Popa, D. and J. Hui, "6LoPLC: Transmission of IPv6 Packets over IEEE 1901.2 Narrowband Powerline Communication Networks", draft-popa-6lo-6loplc-ipv6-over-ieee19012-networks-00 (work in progress), March 2014.
- [RFC3610] Whiting, D., Housley, R., and N. Ferguson, "Counter with CBC-MAC (CCM)", RFC 3610, DOI 10.17487/RFC3610, September 2003, <<http://www.rfc-editor.org/info/rfc3610>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC7102] Vasseur, JP., "Terms Used in Routing for Low-Power and Lossy Networks", RFC 7102, DOI 10.17487/RFC7102, January 2014, <<http://www.rfc-editor.org/info/rfc7102>>.
- [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>>.
- [RFC7668] Nieminen, J., Savolainen, T., Isomaki, M., Patil, B., Shelby, Z., and C. Gomez, "IPv6 over BLUETOOTH(R) Low Energy", RFC 7668, DOI 10.17487/RFC7668, October 2015, <<http://www.rfc-editor.org/info/rfc7668>>.

10.3. External Informative References

- [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. 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.ietf-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 6LN 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 6LN 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 IEEE std 802.15.4 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], IEEE std 802.11ah [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 [RFC7668].

Related requirements are:

Req3.1: The support of the registration mechanism SHOULD be extended to more LLN links than IEEE std 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 IEEE std 802.15.4 [IEEEstd802154] 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 the variation of CCM [RFC3610] called 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.

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

Erik Nordmark
Arista Networks
Santa Clara, CA
USA

Email: nordmark@arista.com

Samita Chakrabarti
San Jose, CA
USA

Email: samitac.ietf@gmail.com

6Lo Working Group
Internet-Draft
Intended status: Informational
Expires: 29 July 2022

Y-G. Hong
Daejeon University
C.G. Gomez
UPC
Y-H. Choi
ETRI
AR. Sangi
Huaiyin Institute of Technology
S. Chakrabarti
January 2022

IPv6 over Constrained Node Networks (6Lo) Applicability & Use cases
draft-ietf-6lo-use-cases-12

Abstract

This document describes the applicability of IPv6 over constrained node networks (6Lo) and provides practical deployment examples. In addition to IEEE Std 802.15.4, various link layer technologies such as ITU-T G.9959 (Z-Wave), Bluetooth Low Energy, DECT-ULE, MS/TP, NFC, and PLC are used as examples. The document targets an audience who would like to understand and evaluate running end-to-end IPv6 over the constrained node networks for local or Internet connectivity.

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 5 July 2022.

Copyright Notice

Copyright (c) 2022 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. 6lo Link layer technologies	4
2.1. ITU-T G.9959	4
2.2. Bluetooth LE	4
2.3. DECT-ULE	5
2.4. MS/TP	5
2.5. NFC	6
2.6. PLC	6
2.7. Comparison between 6lo link layer technologies	8
3. Guidelines for adopting IPv6 stack (6lo)	9
4. 6lo Deployment Scenarios	11
4.1. Wi-SUN usage of 6lo in network layer	11
4.2. Thread usage of 6lo in network layer	13
4.3. G3-PLC usage of 6lo in network layer	13
4.4. Netricity usage of 6lo in network layer	14
5. 6lo Use Case Examples	15
5.1. Use case of ITU-T G.9959: Smart Home	15
5.2. Use case of Bluetooth LE: Smartphone-based Interaction	16
5.3. Use case of DECT-ULE: Smart Home	17
5.4. Use case of MS/TP: Building Automation Networks	17
5.5. Use case of NFC: Alternative Secure Transfer	18
5.6. Use case of PLC: Smart Grid	19
6. IANA Considerations	20
7. Security Considerations	20
8. Acknowledgements	20
9. Informative References	20
Appendix A. Design Space Dimensions for 6lo Deployment	26
Authors' Addresses	28

1. Introduction

Running IPv6 on constrained node networks presents challenges, due to the characteristics of these networks such as small packet size, low power, low bandwidth, low cost, and large number of devices, among others [RFC4919][RFC7228]. For example, many IEEE Std 802.15.4 variants [IEEE802154] exhibit a frame size of 127 octets, whereas IPv6 requires its underlying layer to support an MTU of 1280 bytes.

Furthermore, those IEEE Std 802.15.4 variants do not offer fragmentation and reassembly functionality. (It is noted that IEEE Std 802.15.9-2016 provides multiplexing and fragmentation layer for the IEEE Std 802.15.4[IEEE802159].) Therefore, an appropriate adaptation layer supporting fragmentation and reassembly must be provided below IPv6. Also, the limited IEEE Std 802.15.4 frame size and low energy consumption requirements motivate the need for packet header compression. The IETF IPv6 over Low-Power WPAN (6LoWPAN) working group published a suite of specification that provide an adaptation layer to support IPv6 over IEEE Std 802.15.4 comprising the following functionality:

- * Fragmentation and reassembly, address autoconfiguration, and a frame format [RFC4944],
- * IPv6 (and UDP) header compression [RFC6282],
- * Neighbor Discovery Optimization for 6LoWPAN [RFC6775][RFC8505].

As Internet of Things (IoT) services become more popular, the IETF 6lo working group [IETF_6lo] has defined adaptation layer functionality to support IPv6 over various link layer technologies other than IEEE Std 802.15.4, such as Bluetooth Low Energy (Bluetooth LE), ITU-T G.9959 (Z-Wave), Digital Enhanced Cordless Telecommunications - Ultra Low Energy (DECT-ULE), Master-Slave/Token Passing (MS/TP), Near Field Communication (NFC), and Power Line Communication (PLC). The 6lo adaptation layers use a variation of the 6LoWPAN stack applied to each particular link layer technology.

The 6LoWPAN working group produced the document entitled "Design and Application Spaces for 6LoWPANs" [RFC6568], which describes potential application scenarios and use cases for low-power wireless personal area networks. The present document aims to provide guidance to an audience who are new to the IPv6 over constrained node networks (6lo) concept and want to assess its application to the constrained node network of their interest. This 6lo applicability document describes a few sets of practical 6lo deployment scenarios and use cases examples. In addition, it considers various network design space dimensions such as deployment, network size, power source, connectivity, multi-hop communication, traffic pattern, security level, mobility, and QoS requirements etc.

This document provides the applicability and use cases of 6lo, considering the following aspects:

- * It covers various IoT-related wired/wireless link layer technologies providing practical information of such technologies.

- * It provides a general guideline on how the 6LoWPAN stack can be modified for a given L2 technology.
- * Various 6lo use cases and practical deployment examples are described.

2. 6lo Link layer technologies

2.1. ITU-T G.9959

The ITU-T G.9959 Recommendation [G.9959] targets low-power Wireless Personal Area Networks (WPANs), and defines physical layer and link layer functionality. Physical layers of 9.6 kbit/s, 40 kbit/s and 100 kbit/s are supported. G.9959 defines how a unique 32-bit HomeID network identifier is assigned by a network controller and how an 8-bit NodeID host identifier is allocated to each node. NodeIDs are unique within the network identified by the HomeID. The G.9959 HomeID represents an IPv6 subnet that is identified by one or more IPv6 prefixes [RFC7428]. The ITU-T G.9959 can be used for smart home applications.

2.2. Bluetooth LE

Bluetooth LE was introduced in Bluetooth 4.0, enhanced in Bluetooth 4.1, and developed further in successive versions. Bluetooth SIG has also published the Internet Protocol Support Profile (IPSP). The IPSP enables discovery of IP-enabled devices and establishment of link-layer connection for transporting IPv6 packets. IPv6 over Bluetooth LE is dependent on both Bluetooth 4.1 and IPSP 1.0 or newer.

Many devices such as mobile phones, notebooks, tablets and other handheld computing devices which support Bluetooth 4.0 or subsequent versions also support the low-energy variant of Bluetooth. Bluetooth LE is also being included in many different types of accessories that collaborate with mobile devices. An example of a use case for a Bluetooth LE accessory is a heart rate monitor that sends data via the mobile phone to a server on the Internet [RFC7668]. A typical usage of Bluetooth LE is smartphone-based interaction with constrained devices. Bluetooth LE was originally designed to enable star topology networks. However, recent Bluetooth versions support the formation of extended topologies, and IPv6 support for mesh networks of Bluetooth LE devices is being developed [RFC9159].

2.3. DECT-ULE

DECT-ULE is a low power air interface technology that is designed to support both circuit switched services, such as voice communication, and packet mode data services at modest data rate.

The DECT-ULE protocol stack consists of the physical layer operating at frequencies in the dedicated 1880 - 1920 MHz frequency band depending on the region and uses a symbol rate of 1.152 Mbps. Radio bearers are allocated by use of FDMA/TDMA/TDD techniques.

In its generic network topology, DECT is defined as a cellular network technology. However, the most common configuration is a star network with a single Fixed Part (FP) defining the network with a number of Portable Parts (PP) attached. The Medium Access Control (MAC) layer supports traditional DECT as this is used for services like discovery, pairing, security features etc. All these features have been reused from DECT.

The DECT-ULE device can switch to the ULE mode of operation, utilizing the new ULE MAC layer features. The DECT-ULE Data Link Control (DLC) provides multiplexing as well as segmentation and re-assembly for larger packets from layers above. The DECT-ULE layer also implements per-message authentication and encryption. The DLC layer ensures packet integrity and preserves packet order, but delivery is based on best effort.

The current DECT-ULE MAC layer standard supports low bandwidth data broadcast. However the usage of this broadcast service has not yet been standardized for higher layers [RFC8105]. DECT-ULE can be used for smart metering in a home.

2.4. MS/TP

MS/TP is a MAC protocol for the RS-485 [TIA-485-A] physical layer and is used primarily in building automation networks.

An MS/TP device is typically based on a low-cost microcontroller with limited processing power and memory. These constraints, together with low data rates and a small MAC address space, are similar to those faced in 6LoWPAN networks. MS/TP differs significantly from 6LoWPAN in at least three respects: a) MS/TP devices are typically mains powered, b) all MS/TP devices on a segment can communicate directly so there are no hidden node or mesh routing issues, and c) the latest MS/TP specification provides support for large payloads, eliminating the need for fragmentation and reassembly below IPv6.

MS/TP is designed to enable multidrop networks over shielded twisted pair wiring. It can support network segments up to 1000 meters in length at a data rate of 115.2 kbit/s or segments up to 1200 meters in length at lower bit rates. An MS/TP interface requires only a Universal Asynchronous Receiver-Transmitter (UART), an RS-485 [TIA-485-A] transceiver with a driver that can be disabled, and a 5 ms resolution timer. The MS/TP MAC is typically implemented in software.

Because of its superior "range" (~1 km) compared to many low power wireless data links, MS/TP may be suitable to connect remote devices (such as district heating controllers) to the nearest building control infrastructure over a single link [RFC8163].

2.5. NFC

NFC technology enables simple and safe two-way interactions between electronic devices, allowing consumers to perform contactless transactions, access digital content, and connect electronic devices with a single touch. NFC complements many popular consumer level wireless technologies, by utilizing the key elements in existing standards for contactless card technology (ISO/IEC 14443 A&B and JIS-X 6319-4).

Extending the capability of contactless card technology, NFC also enables devices to share information at a distance that is less than 10 cm with a maximum communication speed of 424 kbps. Users can share business cards, make transactions, access information from a smart poster or provide credentials for access control systems with a simple touch.

NFC's bidirectional communication ability is ideal for establishing connections with other technologies by the simplicity of touch. In addition to the easy connection and quick transactions, simple data sharing is also available [I-D.ietf-6lo-nfc]. NFC can be used for secure transfer in healthcare services.

2.6. PLC

PLC is a data transmission technique that utilizes power conductors as medium [I-D.ietf-6lo-plc]. Unlike other dedicated communication infrastructure, power conductors are widely available indoors and outdoors. Moreover, wired technologies cause less interference to the radio medium than wireless technologies and are more reliable than their wireless counterparts.

The below table shows some available open standards defining PLC.

PLC Systems	Frequency Range	Type	Data Rate	Distance
IEEE1901	<100MHz	Broadband	200Mbps	1000m
IEEE1901.1	<12MHz	PLC-IoT	10Mbps	2000m
IEEE1901.2	<500kHz	Narrowband	200kbps	3000m
G3-PLC	<500kHz	Narrowband	234kbps	3000m

Table 1: Some Available Open Standards in PLC

IEEE Std 1901 [IEEE1901] defines a broadband variant of PLC but is effective within short range. This standard addresses the requirements of applications with high data rate such as: Internet, HDTV, Audio, Gaming etc. Broadband operates on Orthogonal Frequency Division Multiplexing (OFDM) modulation.

IEEE Std 1901.1 [IEEE1901.1] defines a medium frequency band (less than 12 MHz) broadband PLC technology for smart grid applications based on OFDM. By achieving an extended communication range with medium speeds, this standard can be applied both in indoor and outdoor scenarios, such as Advanced Metering Infrastructure (AMI), street lighting, electric vehicle charging, smart city etc.

IEEE Std 1901.2 [IEEE1901.2] defines a narrowband variant of PLC with less data rate but significantly higher transmission range that could be used in an indoor or even an outdoor environment. It is applicable to typical IoT applications such as: Building Automation, Renewable Energy, Advanced Metering, Street Lighting, Electric Vehicle, Smart Grid etc. Moreover, IEEE Std 1901.2 standard is based on the 802.15.4 MAC sub-layer and fully endorses the security scheme defined in 802.15.4 [RFC8036]. A typical use case of PLC is smart grid.

G3-PLC [G3-PLC] is a narrowband PLC technology that is based on the ITU-T G.9903 Recommendation [G.9903]. The ITU-T G.9903 Recommendation contains the physical layer and data link layer specification for the G3-PLC narrowband OFDM power line communication transceivers, for communications via alternating current and direct current electric power lines over frequencies below 500 kHz.

2.7. Comparison between 6lo link layer technologies

In above clauses, various 6lo link layer technologies are described. The following table shows dominant parameters of each use case corresponding to the 6lo link layer technology.

	Z-Wave	BLE	DECT-ULE	MS/TP	NFC	PLC
Usage	Home Auto-mation	Interact w/ Smart Phone	Meter Reading	Building Auto-mation	Health-care Service	Smart Grid
Topology & Subnet	L2-mesh or L3-mesh	Star & Mesh	Star No mesh	MS/TP No mesh	P2P L2-mesh	Star Tree Mesh
Mobility Requirement	No	Low	No	No	Moderate	No
Security Requirement	High + Privacy required	Partially	High + Privacy required	High + Authen. required	High	High + Encrypt. required
Buffering Requirement	Low	Low	Low	Low	Low	Low
Latency, QoS Requirement	High	Low	Low	High	High	Low
Data Rate	Infrequent	Infrequent	Infrequent	Frequent	Small	Infrequent
RFC # or Draft	RFC7428	RFC7668, RFC9159	RFC8105	RFC8163	draft-ietf-6lo-nfc	draft-ietf-6lo-plc

Table 2: Comparison between 6lo link layer technologies

3. Guidelines for adopting IPv6 stack (6lo)

6lo aims at reusing and/or adapting existing 6LoWPAN functionality in order to efficiently support IPv6 over a variety of IoT L2 technologies. The following guideline targets new candidate constrained L2 technologies that may be considered for running a modified 6LoWPAN stack on top. The modification of 6LoWPAN stack should be based on the following:

- * **Addressing Model:** Addressing model determines whether the device is capable of forming IPv6 link-local and global addresses and what is the best way to derive the IPv6 addresses for the constrained L2 devices. L2-address-derived IPv6 addresses are specified in [RFC4944], but there exist implications for privacy. For global usage, a unique IPv6 address must be derived using an assigned prefix and a unique interface ID. [RFC8065] provides such guidelines. For MAC-derived IPv6 addresses, please refer to [RFC8163] for IPv6 address mapping examples. Broadcast and multicast support are dependent on the L2 networks. Most low-power L2 implementations map multicast to broadcast networks. So care must be taken in the design when to use broadcast and try to stick to unicast messaging whenever possible.
- * **MTU Considerations:** The deployment should consider packet maximum transmission unit (MTU) needs over the link layer and should consider if fragmentation and reassembly of packets are needed at the 6LoWPAN layer. For example, if the link layer supports fragmentation and reassembly of packets, then the 6LoWPAN layer may not need to support fragmentation/reassembly. In fact, for most efficiency, choosing a low-power link layer that can carry unfragmented application packets would be optimum for packet transmission if the deployment can afford it. Please refer to 6lo RFCs [RFC7668], [RFC8163], [RFC8105] for example guidance.
- * **Mesh or L3-Routing:** 6LoWPAN specifications provide mechanisms to support mesh routing at L2, a configuration called mesh-under [RFC6606]. It is also possible to use an L3 routing protocol in 6LoWPAN, an approach known as route-over. [RFC6550] defines RPL, a L3 routing protocol for low power and lossy networks using directed acyclic graphs. 6LoWPAN is routing-protocol-agnostic and does not specify any particular L2 or L3 routing protocol to use with a 6LoWPAN stack.
- * **Address Assignment:** 6LoWPAN developed a new version of IPv6 Neighbor Discovery [RFC4861][RFC4862]. 6LoWPAN Neighbor Discovery [RFC6775][RFC8505] inherits from IPv6 Neighbor Discovery for mechanisms such as Stateless Address Autoconfiguration (SLAAC) and Neighbor Unreachability Detection (NUD). A 6LoWPAN node is also

expected to be an IPv6 host per [RFC8200] which means it should ignore consumed routing headers and Hop-by-Hop options; when operating in a RPL network [RFC6550], it is also beneficial to support IP-in-IP encapsulation [RFC9008]. The 6LoWPAN node should also support [RFC8505] and use it as the default Neighbor Discovery method. It is the responsibility of the deployment to ensure unique global IPv6 addresses for Internet connectivity. For local-only connectivity IPv6 Unique Local Address (ULA) may be used. [RFC6775][RFC8505] specifies the 6LoWPAN border router (6LBR), which is responsible for prefix assignment to the 6LoWPAN network. A 6LBR can be connected to the Internet or to an enterprise network via one of the interfaces. Please refer to [RFC7668] and [RFC8105] for examples of address assignment considerations. In addition, privacy considerations [RFC8065] must be consulted for applicability. In certain scenarios, the deployment may not support IPv6 address autoconfiguration due to regulatory and business reasons and may choose to offer a separate address assignment service. Address Protection for 6LoWPAN Neighbor Discovery (AP-ND) [RFC8928] enables Source Address Validation [RFC6620] and protects the address ownership against impersonation attacks.

- * **Broadcast Avoidance:** 6LoWPAN Neighbor Discovery aims at reducing the amount of multicast traffic of classical Neighbor Discovery, since IP-level multicast translates into L2 broadcast in many L2 technologies. 6LoWPAN Neighbor Discovery relies on a proactive registration to avoid the use of multicast for address resolution. It also uses a unicast method for Duplicate Address Detection (DAD), and avoids multicast lookups from all nodes by using non-onlink prefixes. Router Advertisements (RAs) are also sent in unicast, in response to Router Solicitations (RSs)
- * **Host-to-Router interface:** 6lo has defined registration extensions for 6LoWPAN Neighbor Discovery [RFC8505]. This effort provides a host-to-router interface by which a host can request its router to ensure reachability for the address registered with the router. Note that functionality has been developed to ensure that such a host can benefit from routing services in a RPL network [RFC9010]
- * **Proxy Neighbor Discovery:** Further functionality also allows a device (e.g. an energy-constrained device that needs to sleep most of the time) to request proxy Neighbor Discovery services from a 6LoWPAN Backbone Router (6BBR) [RFC8505][RFC8929]. The latter federates a number of links into a multilink subnet.
- * **Header Compression:** IPv6 header compression [RFC6282] is a vital part of IPv6 over low power communication. Examples of header compression over different link-layer specifications are found in

[RFC7668], [RFC8163], [RFC8105]. A generic header compression technique is specified in [RFC7400]. For 6LoWPAN networks where RPL is the routing protocol, there exist 6LoWPAN header compression extensions which allow to compress also the RPL artifacts used when forwarding packets in the route-over mesh [RFC8138] [RFC9035]

- * Security and Encryption: Though 6LoWPAN basic specifications do not address security at the network layer, the assumption is that L2 security must be present. In addition, application-level security is highly desirable. The working groups [IETF_ace] and [IETF_core] should be consulted for application and transport level security. 6lo working group is working on address authentication [RFC8928] and secure bootstrapping is also being discussed at IETF. However, there may be different levels of security available in a deployment through other standards such as hardware-level security or certificates for initial booting process. Encryption is important if the implementation can afford it.
- * Additional processing: [RFC8066] defines guidelines for ESC dispatch octets use in the 6LoWPAN header. An implementation may take advantage of ESC header to offer a deployment specific processing of 6LoWPAN packets.

4. 6lo Deployment Scenarios

4.1. Wi-SUN usage of 6lo in network layer

Wireless Smart Ubiquitous Network (Wi-SUN) [Wi-SUN] is a technology based on the IEEE Std 802.15.4g standard. Wi-SUN networks support star and mesh topologies, as well as hybrid star/mesh deployments, but these are typically laid out in a mesh topology where each node relays data for the network to provide network connectivity. Wi-SUN networks are deployed on both powered and battery-operated devices [RFC8376].

The main application domains targeted by Wi-SUN are smart utility and smart city networks. This includes, but is not limited to the following applications:

- * Advanced Metering Infrastructure
- * Distribution Automation
- * Home Energy Management
- * Infrastructure Management

- * Intelligent Transportation Systems
- * Smart Street Lighting
- * Agriculture
- * Structural health (bridges, buildings)
- * Monitoring and Asset Management
- * Smart Thermostats, Air Conditioning and Heat Controls
- * Energy Usage Information Displays

The Wi-SUN Alliance Field Area Network (FAN) covers primarily outdoor networks, and its specification is oriented towards meeting the more rigorous challenges of these environments. It has the following features:

- * Open standards based on IEEE802, IETF, TIA, ETSI
- * Architecture based on an IPv6 frequency hopping wireless mesh network with enterprise-level security
- * Simple infrastructure of low cost, low complexity
- * Enhanced network robustness, reliability, and resilience to interference, due to high redundancy and frequency hopping
- * Enhanced scalability, long range, and energy friendliness
- * Supports multiple global license-exempt sub-GHz bands
- * Multi-vendor interoperability
- * Very low power modes in development permitting long term battery operation of network nodes

The Wi-SUN FAN specification defines an IPv6-based protocol suite including TCP/UDP, IPv6, 6lo adaptation layer, DHCPv6 for IPv6 address management, RPL, and ICMPv6.

4.2. Thread usage of 6lo in network layer

Thread is an IPv6-based networking protocol stack built on open standards, designed for smart home environments, and based on low-power IEEE Std 802.15.4 mesh networks. Because of its IPv6 foundation, Thread can support existing popular application layers and IoT platforms, provide end-to-end security, ease development and enable flexible and future-proof designs [Thread].

The Thread specification uses the IEEE Std 802.15.4 [IEEE802154] physical and MAC layers operating at 250 kbps in the 2.4 GHz band.

Thread devices use 6LoWPAN, as defined in [RFC4944][RFC6282], for transmission of IPv6 Packets over IEEE Std 802.15.4 networks. Header compression is used within the Thread network and devices transmitting messages compress the IPv6 header to minimize the size of the transmitted packet. The mesh header is supported for link-layer (i.e., mesh under) forwarding. The mesh header as used in Thread also allows efficient end-to-end fragmentation of messages rather than the hop-by-hop fragmentation specified in [RFC4944]. Mesh under routing in Thread is based on a distance vector protocol in a full mesh topology.

4.3. G3-PLC usage of 6lo in network layer

G3-PLC [G3-PLC] is a narrowband PLC technology that is based on the ITU-T G.9903 Recommendation [G.9903]. G3-PLC supports multi-hop mesh network topology, and facilitates highly-reliable, long-range communication. With the abilities to support IPv6 and to cross transformers, G3-PLC is regarded as one of the next-generation narrowband PLC technologies. G3-PLC has got massive deployments over several countries, e.g. Japan and France.

The main application domains targeted by G3-PLC are smart grid and smart cities. This includes, but is not limited to the following applications:

- * Smart Metering
- * Vehicle-to-Grid Communication
- * Demand Response
- * Distribution Automation
- * Home/Building Energy Management Systems
- * Smart Street Lighting

- * Advanced Metering Infrastructure (AMI) backbone network
- * Wind/Solar Farm Monitoring

In the G3-PLC specification, the 6lo adaption layer utilizes the 6LoWPAN functions (e.g. header compression, fragmentation and reassembly). However, due to the different characteristics of the PLC media, the 6LoWPAN adaptation layer cannot perfectly fulfill the requirements [I-D.ietf-6lo-plc]. The ESC dispatch type is used in the G3-PLC to provide native mesh routing and bootstrapping functionalities [RFC8066].

4.4. Netricity usage of 6lo in network layer

The Netricity program in HomePlug Powerline Alliance [NETRICITY] promotes the adoption of products built on the IEEE Std 1901.2 low-frequency narrowband PLC standard, which provides for urban and long distance communications and propagation through transformers of the distribution network using frequencies below 500 kHz. The technology also addresses requirements that assure communication privacy and secure networks.

The main application domains targeted by Netricity are smart grid and smart cities. This includes, but is not limited to the following applications:

- * Utility grid modernization
- * Distribution automation
- * Meter-to-Grid connectivity
- * Micro-grids
- * Grid sensor communications
- * Load control
- * Demand response
- * Net metering
- * Street Lighting control
- * Photovoltaic panel monitoring

Netricity system architecture is based on the physical and MAC layers of IEEE Std 1901.2 PLC standard. Regarding the 6lo adaptation layer and IPv6 network layer, Netricity utilizes IPv6 protocol suite including 6lo/6LoWPAN header compression, DHCPv6 for IP address management, RPL routing protocol, ICMPv6, and unicast/multicast forwarding. Note that the L3 routing in Netricity uses RPL in non-storing mode with the MRHOF objective function based on the own defined Estimated Transmission Time (ETT) metric.

5. 6lo Use Case Examples

As IPv6 stacks for constrained node networks use a variation of the 6LoWPAN stack applied to each particular link layer technology, various 6lo use cases can be provided. In this section, various 6lo use cases which are based on different link layer technologies are described.

5.1. Use case of ITU-T G.9959: Smart Home

Z-Wave is one of the main technologies that may be used to enable smart home applications. Born as a proprietary technology, Z-Wave was specifically designed for this particular use case. Recently, the Z-Wave radio interface (physical and MAC layers) has been standardized as the ITU-T G.9959 specification.

Example: Use of ITU-T G.9959 for Home Automation

Variety of home devices (e.g. light dimmers/switches, plugs, thermostats, blinds/curtains and remote controls) are augmented with ITU-T G.9959 interfaces. A user may turn on/off or may control home appliances by pressing a wall switch or by pressing a button in a remote control. Scenes may be programmed, so that after a given event, the home devices adopt a specific configuration. Sensors may also periodically send measurements of several parameters (e.g. gas presence, light, temperature, humidity, etc.) which are collected at a sink device, or may generate commands for actuators (e.g. a smoke sensor may send an alarm message to a safety system).

The devices involved in the described scenario are nodes of a network that follows the mesh topology, which is suitable for path diversity to face indoor multipath propagation issues. The multihop paradigm allows end-to-end connectivity when direct range communication is not possible. Security support is required, specially for safety-related communication. When a user interaction (e.g. a button press) triggers a message that encapsulates a command, if the message is lost, the user may have to perform further interactions to achieve the desired effect (e.g. turning off a light). A reaction to a user interaction will be perceived by the user as immediate as long as the reaction takes place within 0.5 seconds [RFC5826].

5.2. Use case of Bluetooth LE: Smartphone-based Interaction

The key feature behind the current high Bluetooth LE momentum is its support in a large majority of smartphones in the market. Bluetooth LE can be used to allow the interaction between the smartphone and surrounding sensors or actuators. Furthermore, Bluetooth LE is also the main radio interface currently available in wearables. Since a smartphone typically has several radio interfaces that provide Internet access, such as Wi-Fi or 4G, the smartphone can act as a gateway for nearby devices such as sensors, actuators or wearables. Bluetooth LE may be used in several domains, including healthcare, sports/wellness and home automation.

Example: Use of Bluetooth LE-based Body Area Network for fitness

A person wears a smartwatch for fitness purposes. The smartwatch has several sensors (e.g. heart rate, accelerometer, gyrometer, GPS, temperature, etc.), a display, and a Bluetooth LE radio interface. The smartwatch can show fitness-related statistics on its display. However, when a paired smartphone is in the range of the smartwatch, the latter can report almost real-time measurements of its sensors to the smartphone, which can forward the data to a cloud service on the Internet. 6lo enables this use case by providing efficient end-to-end IPv6 support. In addition, the smartwatch can receive notifications (e.g. alarm signals) from the cloud service via the smartphone. On the other hand, the smartphone may locally generate messages for the smartwatch, such as e-mail reception or calendar notifications.

The functionality supported by the smartwatch may be complemented by other devices such as other on-body sensors, wireless headsets or head-mounted displays. All such devices may connect to the smartphone creating a star topology network whereby the smartphone is the central component. Support for extended network topologies (e.g. mesh networks) is being developed as of the writing.

5.3. Use case of DECT-ULE: Smart Home

DECT is a technology widely used for wireless telephone communications in residential scenarios. Since DECT-ULE is a low-power variant of DECT, DECT-ULE can be used to connect constrained devices such as sensors and actuators to a Fixed Part, a device that typically acts as a base station for wireless telephones. Therefore, DECT-ULE is specially suitable for the connected home space in application areas such as home automation, smart metering, safety, healthcare, etc. Since DECT-ULE uses dedicated bandwidth, it avoids the coexistence issues suffered by other technologies that use e.g. ISM frequency bands.

Example: Use of DECT-ULE for Smart Metering

The smart electricity meter of a home is equipped with a DECT-ULE transceiver. This device is in the coverage range of the Fixed Part of the home. The Fixed Part can act as a router connected to the Internet. This way, the smart meter can transmit electricity consumption readings through the DECT-ULE link with the Fixed Part, and the latter can forward such readings to the utility company using Wide Area Network (WAN) links. The meter can also receive queries from the utility company or from an advanced energy control system controlled by the user, which may also be connected to the Fixed Part via DECT-ULE.

5.4. Use case of MS/TP: Building Automation Networks

The primary use case for IPv6 over MS/TP (6LoBAC) is in building automation networks. [BACnet] is the open international standard protocol for building automation, and MS/TP is defined in [BACnet] Clause 9. MS/TP was designed to be a low cost multi-drop field bus to inter-connect the most numerous elements (sensors and actuators) of a building automation network to their controllers. A key aspect of 6LoBAC is that it is designed to co-exist with BACnet MS/TP on the same link, easing the ultimate transition of some BACnet networks to native end-to-end IPv6 transport protocols. New applications for 6LoBAC may be found in other domains where low cost, long distance, and low latency are required. Note that BACnet comprises various networking solutions other than MS/TP, including the recently emerged BACnet IP. However, the latter is based on high speed Ethernet infrastructure, and thus it falls outside of the constrained node network scope.

Example: Use of 6LoBAC in Building Automation Networks

The majority of installations for MS/TP are for "terminal" or "unitary" controllers, i.e. single zone or room controllers that may connect to HVAC or other controls such as lighting or blinds. The economics of daisy-chaining a single twisted-pair between multiple devices is often preferred over home-run Cat-5 style wiring.

A multi-zone controller might be implemented as an IP router between a traditional Ethernet link and several 6LoBAC links, fanning out to multiple terminal controllers.

The superior distance capabilities of MS/TP (~1 km) compared to other 6Lo media may suggest its use in applications to connect remote devices to the nearest building infrastructure. For example, remote pumping or measuring stations with moderate bandwidth requirements can benefit from the low cost and robust capabilities of MS/TP over other wired technologies such as DSL, and without the line-of-sight restrictions or hop-by-hop latency of many low cost wireless solutions.

5.5. Use case of NFC: Alternative Secure Transfer

In different applications, a variety of secured data can be handled and transferred. Depending on the security level of the data, different transfer methods can be alternatively selected.

Example: Use of NFC for Secure Transfer in Healthcare Services with Tele-Assistance

A senior citizen who lives alone wears one to several wearable 6Lo devices to measure heartbeat, pulse rate, etc. The 6Lo devices are densely installed at home for movement detection. A 6LBR at home will send the sensed information to a connected healthcare center. Portable base stations with LCDs may be used to check the data at home, as well. Data is gathered in both periodic and event-driven fashion. In this application, event-driven data can be very time-critical. In addition, privacy also becomes a serious issue in this case, as the sensed data is very personal.

While the senior citizen is provided audio and video healthcare services by a tele-assistance based on LTE connections, the senior citizen can alternatively use NFC connections to transfer the personal sensed data to the tele-assistance. Hidden hackers can overhear the data based on the LTE connection, but they cannot gather the personal data over the NFC connection.

5.6. Use case of PLC: Smart Grid

The smart grid concept is based on deploying numerous operational and energy measuring sub-systems in an electricity grid system. It comprises multiple administrative levels/segments to provide connectivity among these numerous components. Last mile connectivity is established over the Low Voltage (LV) segment, whereas connectivity over electricity distribution takes place in the High Voltage (HV) segment. Smart grid systems include Advanced Metering Infrastructure (AMI), Demand Response (DR), Home Energy Management System (HEMS), Wide Area Situational Awareness (WASA), among others.

Although other wired and wireless technologies are also used in Smart Grid, PLC enjoys the advantage of reliable data communication over electrical power lines that are already present, and the deployment cost can be comparable to wireless technologies. The 6lo-related scenarios for PLC mainly lie in the LV PLC networks with most applications in the area of Advanced Metering Infrastructure, Vehicle-to-Grid communications, in-home energy management and smart street lighting.

Example: Use of PLC for Advanced Metering Infrastructure

Household electricity meters transmit time-based data of electric power consumption through PLC. Data concentrators receive all the meter data in their corresponding living districts and send them to the Meter Data Management System (MDMS) through WAN network (e.g. Medium-Voltage PLC, Ethernet or GPRS) for storage and analysis. Two-way communications are enabled which means smart meters can do actions like notification of electricity charges according to the commands from the utility company.

With the existing power line infrastructure as communication medium, cost on building up the PLC network is naturally saved, and more importantly, labor operational costs can be minimized from a long-term perspective. Furthermore, this AMI application speeds up electricity charge, reduces losses by restraining power theft and helps to manage the health of the grid based on line loss analysis.

Example: Use of PLC (IEEE Std 1901.1) for WASA in Smart Grid

Many sub-systems of Smart Grid require low data rate and narrowband variants (e.g., IEEE Std 1901.1) of PLC fulfill such requirements. Recently, more complex scenarios are emerging that require higher data rates.

WASA sub-system is an appropriate example that collects large amount of information about the current state of the grid over wide area from electric substations as well as power transmission lines. The collected feedback is used for monitoring, controlling and protecting all the sub-systems.

6. IANA Considerations

There are no IANA considerations related to this document.

7. Security Considerations

Security considerations are not directly applicable to this document. For the use cases, the security requirements described in the protocol specifications apply.

8. Acknowledgements

Carles Gomez has been funded in part by the Spanish Government through the Jose Castillejo CAS15/00336 grant, the TEC2016-79988-P grant, and the PID2019-106808RA-I00 grant, and by Secretaria d'Universitats i Recerca del Departament d'Empresa i Coneixement de la Generalitat de Catalunya 2017 through grant SGR 376. His contribution to this work has been carried out in part during his stay as a visiting scholar at the Computer Laboratory of the University of Cambridge.

Thomas Watteyne, Pascal Thubert, Xavier Vilajosana, Daniel Migault, Jianqiang Hou, Kerry Lynn, S.V.R. Anand, and Seyed Mahdi Darroudi have provided valuable feedback for this draft.

Das Subir and Michel Veillette have provided valuable information of jupiterMesh and Paul Duffy has provided valuable information of Wi-SUN for this draft. Also, Jianqiang Hou has provided valuable information of G3-PLC and Netricity for this draft. Take Aanstoot, Kerry Lynn, and Dave Robin have provided valuable information of MS/TP and practical use case of MS/TP for this draft.

Deoknyong Ko has provided relevant text of LTE-MTC and he shared his experience to deploy IPv6 and 6lo technologies over LTE MTC in SK Telecom.

9. Informative References

- [BACnet] "ASHRAE, "BACnet-A Data Communication Protocol for Building Automation and Control Networks", ANSI/ASHRAE Standard 135-2016", January 2016,
<http://www.techstreet.com/ashrae/standards/ashrae-135-2016?product_id=1918140#jumps>.
- [G.9903] "International Telecommunication Union, "Narrowband orthogonal frequency division multiplexing power line communication transceivers for G3-PLC networks", ITU-T Recommendation", August 2017.
- [G.9959] "International Telecommunication Union, "Short range narrow-band digital radiocommunication transceivers - PHY and MAC layer specifications", ITU-T Recommendation", January 2015.
- [G3-PLC] "G3-PLC Alliance", <<http://www.g3-plc.com/home/>>.
- [IEEE1901] "IEEE Standard, IEEE Std 1901-2010 - IEEE Standard for Broadband over Power Line Networks: Medium Access Control and Physical Layer Specifications", 2010,
<<https://standards.ieee.org/findstds/standard/1901-2010.html>>.
- [IEEE1901.1]
"IEEE Standard, IEEE Std 1901.1-2018 - IEEE Standard for Medium Frequency (less than 12 MHz) Power Line Communications for Smart Grid Applications", 2018,
<<https://ieeexplore.ieee.org/document/8360785>>.
- [IEEE1901.2]
"IEEE Standard, IEEE Std 1901.2-2013 - IEEE Standard for Low-Frequency (less than 500 kHz) Narrowband Power Line Communications for Smart Grid Applications", 2013,
<<https://standards.ieee.org/findstds/standard/1901.2-2013.html>>.
- [IEEE802154]
IEEE standard for Information Technology, "IEEE Standard for Low-Rate Wireless Networks".
- [IEEE802159]
IEEE standard for Information Technology, "IEEE Std 802.15.9-2016 - IEEE Recommended Practice for Transport of Key Management Protocol (KMP) Datagrams".

- [I-D.ietf-6lo-nfc]
Choi, Y., Hong, Y., Youn, J., Kim, D., and J. Choi,
"Transmission of IPv6 Packets over Near Field
Communication", Work in Progress, Internet-Draft, draft-
ietf-6lo-nfc-17, 23 August 2020, <[https://www.ietf.org/
internet-drafts/draft-ietf-6lo-nfc-17.txt](https://www.ietf.org/internet-drafts/draft-ietf-6lo-nfc-17.txt)>.
- [I-D.ietf-6lo-plc]
Hou, J., Liu, B. R., Hong, Y., Tang, X., and C. E.
Perkins, "Transmission of IPv6 Packets over PLC Networks",
Work in Progress, Internet-Draft, draft-ietf-6lo-plc-09,
10 January 2022, <[https://www.ietf.org/archive/id/draft-
ietf-6lo-plc-09.txt](https://www.ietf.org/archive/id/draft-ietf-6lo-plc-09.txt)>.
- [IETF_6lo] "IETF IPv6 over Networks of Resource-constrained Nodes
(6lo) working group",
<<https://datatracker.ietf.org/wg/6lo/charter/>>.
- [IETF_ace] "IETF Authentication and Authorization for Constrained
Environments (ace) working group",
<<https://datatracker.ietf.org/wg/ace/charter/>>.
- [IETF_core]
"IETF Constrained RESTful Environments (core) working
group", <<https://datatracker.ietf.org/wg/core/charter/>>.
- [Wi-SUN] "Wi-SUN Alliance", <<http://www.wi-sun.org>>.
- [Thread] "Thread Group", <<https://www.threadgroup.org/Support>>.
- [NETRICITY]
"Netricity program in HomePlug Powerline Alliance",
<<http://groups.homeplug.org/tech/Netricity>>.
- [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>>.

- [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>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<https://www.rfc-editor.org/info/rfc4944>>.
- [RFC5826] Brandt, A., Buron, J., and G. Porcu, "Home Automation Routing Requirements in Low-Power and Lossy Networks", RFC 5826, DOI 10.17487/RFC5826, April 2010, <<https://www.rfc-editor.org/info/rfc5826>>.
- [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, <<https://www.rfc-editor.org/info/rfc6282>>.
- [RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <<https://www.rfc-editor.org/info/rfc6550>>.
- [RFC6568] Kim, E., Kaspar, D., and JP. Vasseur, "Design and Application Spaces for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6568, DOI 10.17487/RFC6568, April 2012, <<https://www.rfc-editor.org/info/rfc6568>>.
- [RFC6606] Kim, E., Kaspar, D., Gomez, C., and C. Bormann, "Problem Statement and Requirements for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing", RFC 6606, DOI 10.17487/RFC6606, May 2012, <<https://www.rfc-editor.org/info/rfc6606>>.
- [RFC6620] Nordmark, E., Bagnulo, M., and E. Levy-Abegnoli, "FCFS SAVI: First-Come, First-Served Source Address Validation Improvement for Locally Assigned IPv6 Addresses", RFC 6620, DOI 10.17487/RFC6620, May 2012, <<https://www.rfc-editor.org/info/rfc6620>>.

- [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>>.
- [RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <<https://www.rfc-editor.org/info/rfc7228>>.
- [RFC7400] Bormann, C., "6LoWPAN-GHC: Generic Header Compression for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 7400, DOI 10.17487/RFC7400, November 2014, <<https://www.rfc-editor.org/info/rfc7400>>.
- [RFC7428] Brandt, A. and J. Buron, "Transmission of IPv6 Packets over ITU-T G.9959 Networks", RFC 7428, DOI 10.17487/RFC7428, February 2015, <<https://www.rfc-editor.org/info/rfc7428>>.
- [RFC7668] Nieminen, J., Savolainen, T., Isomaki, M., Patil, B., Shelby, Z., and C. Gomez, "IPv6 over BLUETOOTH(R) Low Energy", RFC 7668, DOI 10.17487/RFC7668, October 2015, <<https://www.rfc-editor.org/info/rfc7668>>.
- [RFC8036] Cam-Winget, N., Ed., Hui, J., and D. Popa, "Applicability Statement for the Routing Protocol for Low-Power and Lossy Networks (RPL) in Advanced Metering Infrastructure (AMI) Networks", RFC 8036, DOI 10.17487/RFC8036, January 2017, <<https://www.rfc-editor.org/info/rfc8036>>.
- [RFC8065] Thaler, D., "Privacy Considerations for IPv6 Adaptation-Layer Mechanisms", RFC 8065, DOI 10.17487/RFC8065, February 2017, <<https://www.rfc-editor.org/info/rfc8065>>.
- [RFC8066] Chakrabarti, S., Montenegro, G., Droms, R., and J. Woodyatt, "IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) ESC Dispatch Code Points and Guidelines", RFC 8066, DOI 10.17487/RFC8066, February 2017, <<https://www.rfc-editor.org/info/rfc8066>>.
- [RFC8105] Mariager, P., Petersen, J., Ed., Shelby, Z., Van de Logt, M., and D. Barthel, "Transmission of IPv6 Packets over Digital Enhanced Cordless Telecommunications (DECT) Ultra Low Energy (ULE)", RFC 8105, DOI 10.17487/RFC8105, May 2017, <<https://www.rfc-editor.org/info/rfc8105>>.

- [RFC8138] Thubert, P., Ed., Bormann, C., Toutain, L., and R. Cragie, "IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing Header", RFC 8138, DOI 10.17487/RFC8138, April 2017, <<https://www.rfc-editor.org/info/rfc8138>>.
- [RFC8163] Lynn, K., Ed., Martocci, J., Neilson, C., and S. Donaldson, "Transmission of IPv6 over Master-Slave/Token-Passing (MS/TP) Networks", RFC 8163, DOI 10.17487/RFC8163, May 2017, <<https://www.rfc-editor.org/info/rfc8163>>.
- [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>>.
- [RFC8352] Gomez, C., Kovatsch, M., Tian, H., and Z. Cao, Ed., "Energy-Efficient Features of Internet of Things Protocols", RFC 8352, DOI 10.17487/RFC8352, April 2018, <<https://www.rfc-editor.org/info/rfc8352>>.
- [RFC8376] Farrell, S., Ed., "Low-Power Wide Area Network (LPWAN) Overview", RFC 8376, DOI 10.17487/RFC8376, May 2018, <<https://www.rfc-editor.org/info/rfc8376>>.
- [RFC8505] Thubert, P., Ed., Nordmark, E., Chakrabarti, S., and C. Perkins, "Registration Extensions for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Neighbor Discovery", RFC 8505, DOI 10.17487/RFC8505, November 2018, <<https://www.rfc-editor.org/info/rfc8505>>.
- [RFC8928] Thubert, P., Ed., Sarikaya, B., Sethi, M., and R. Struik, "Address-Protected Neighbor Discovery for Low-Power and Lossy Networks", RFC 8928, DOI 10.17487/RFC8928, November 2020, <<https://www.rfc-editor.org/info/rfc8928>>.
- [RFC8929] Thubert, P., Ed., Perkins, C.E., and E. Levy-Abegnoli, "IPv6 Backbone Router", RFC 8929, DOI 10.17487/RFC8929, November 2020, <<https://www.rfc-editor.org/info/rfc8929>>.
- [RFC9008] Robles, M.I., Richardson, M., and P. Thubert, "Using RPI Option Type, Routing Header for Source Routes, and IPv6-in-IPv6 Encapsulation in the RPL Data Plane", RFC 9008, DOI 10.17487/RFC9008, April 2021, <<https://www.rfc-editor.org/info/rfc9008>>.

- [RFC9010] Thubert, P., Ed. and M. Richardson, "Routing for RPL (Routing Protocol for Low-Power and Lossy Networks) Leaves", RFC 9010, DOI 10.17487/RFC9010, April 2021, <<https://www.rfc-editor.org/info/rfc9010>>.
- [RFC9035] Thubert, P., Ed. and L. Zhao, "A Routing Protocol for Low-Power and Lossy Networks (RPL) Destination-Oriented Directed Acyclic Graph (DODAG) Configuration Option for the 6LoWPAN Routing Header", RFC 9035, DOI 10.17487/RFC9035, April 2021, <<https://www.rfc-editor.org/info/rfc9035>>.
- [RFC9159] Gomez, C., Darroudi, S.M., Savolainen, T., and M. Spoerk, "IPv6 Mesh over BLUETOOTH(R) Low Energy Using the Internet Protocol Support Profile (IPSP)", RFC 9159, DOI 10.17487/RFC9159, December 2021, <<https://www.rfc-editor.org/info/rfc9159>>.
- [TIA-485-A] "TIA, "Electrical Characteristics of Generators and Receivers for Use in Balanced Digital Multipoint Systems", TIA-485-A (Revision of TIA-485)", March 2003, <https://global.ihs.com/doc_detail.cfm?item_s_key=00032964>.

Appendix A. Design Space Dimensions for 6lo Deployment

The [RFC6568] lists the dimensions used to describe the design space of wireless sensor networks in the context of the 6LoWPAN working group. The design space is already limited by the unique characteristics of a LoWPAN (e.g. low power, short range, low bit rate). In [RFC6568], the following design space dimensions are described: Deployment, Network size, Power source, Connectivity, Multi-hop communication, Traffic pattern, Mobility, Quality of Service (QoS). However, in this document, the following design space dimensions are considered:

- * Deployment/Bootstrapping: 6lo nodes can be connected randomly, or in an organized manner. The bootstrapping has different characteristics for each link layer technology.
- * Topology: Topology of 6lo networks may inherently follow the characteristics of each link layer technology. Point-to-point, star, tree or mesh topologies can be configured, depending on the link layer technology considered.

- * L2-Mesh or L3-Mesh: L2-mesh and L3-mesh may inherently follow the characteristics of each link layer technology. Some link layer technologies may support L2-mesh and some may not support.
- * Multi-link subnet, single subnet: The selection of multi-link subnet and single subnet depends on connectivity and the number of 6lo nodes.
- * Data rate: Typically, the link layer technologies of 6lo have low rate of data transmission. But, by adjusting the MTU, it can deliver higher upper layer data rate.
- * Buffering requirements: Some 6lo use case may require more data rate than the link layer technology support. In this case, a buffering mechanism to manage the data is required.
- * Security and Privacy Requirements: Some 6lo use case can involve transferring some important and personal data between 6lo nodes. In this case, high-level security support is required.
- * Mobility across 6lo networks and subnets: The movement of 6lo nodes depends on the 6lo use case. If the 6lo nodes can move or moved around, a mobility management mechanism is required.
- * Time synchronization requirements: The requirement of time synchronization of the upper layer service is dependent on the 6lo use case. For some 6lo use case related to health service, the measured data must be recorded with exact time and must be transferred with time synchronization.
- * Reliability and QoS: Some 6lo use case requires high reliability, for example real-time service or health-related services.
- * Traffic patterns: 6lo use cases may involve various traffic patterns. For example, some 6lo use case may require short data length and random transmission. Some 6lo use case may require continuous data and periodic data transmission.
- * Security Bootstrapping: Without the external operations, 6lo nodes must have the security bootstrapping mechanism.
- * Power use strategy: to enable certain use cases, there may be requirements on the class of energy availability and the strategy followed for using power for communication [RFC7228]. Each link layer technology defines a particular power use strategy which may be tuned [RFC8352]. Readers are expected to be familiar with [RFC7228] terminology.

- * Update firmware requirements: Most 6lo use cases will need a mechanism for updating firmware. In these cases support for over the air updates are required, probably in a broadcast mode when bandwidth is low and the number of identical devices is high.
- * Wired vs. Wireless: Plenty of 6lo link layer technologies are wireless, except MS/TP and PLC. The selection of wired or wireless link layer technology is mainly dependent on the requirement of 6lo use cases and the characteristics of wired/wireless technologies. For example, some 6lo use cases may require easy and quick deployment, whereas others may need a continuous source of power.

Authors' Addresses

Yong-Geun Hong
Daejeon University
62 Daehak-ro, Dong-gu
Daejeon

Phone: +82 42 280 4841
Email: yonggeun.hong@gmail.com

Carles Gomez
Universitat Politecnica de Catalunya/Fundacio i2cat
C/Esteve Terradas, 7
08860 Castelldefels
Spain

Email: carlesgo@entel.upc.edu

Younghwan Choi
ETRI
218 Gajeongno, Yuseong
Daejeon

Phone: +82 42 860 1429
Email: yhc@etri.re.kr

Abdur Rashid Sangi
Huaiyin Institute of Technology
No.89 North Beijing Road, Qinghe District
Huaian
223001
P.R. China

Email: sangi_bahrian@yahoo.com

Samita Chakrabarti
San Jose, CA,
United States of America

Email: samitac.ietf@gmail.com

6lo
Internet-Draft
Intended status: Standards Track
Expires: July 14, 2017

P. Thubert, Ed.
Cisco Systems
J. Hui
Nest Labs
January 10, 2017

LLN Fragment Forwarding and Recovery
draft-thubert-6lo-forwarding-fragments-04

Abstract

In order to be routed, a fragmented 6LoWPAN packet must be reassembled at every hop of a multihop link where lower layer fragmentation occurs. Considering that the IPv6 minimum MTU is 1280 bytes and that an 802.15.4 frame can have a payload limited to 74 bytes in the worst case, a packet might end up fragmented into as many as 18 fragments at the 6LoWPAN shim layer. If a single one of those fragments is lost in transmission, all fragments must be resent, further contributing to the congestion that might have caused the initial packet loss. This draft introduces a simple protocol to forward and recover individual fragments that might be lost over multiple hops between 6LoWPAN endpoints.

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 July 14, 2017.

Copyright Notice

Copyright (c) 2017 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. Terminology	3
3. Rationale	4
4. Requirements	5
5. Overview	6
6. New Dispatch types and headers	8
6.1. Recoverable Fragment Dispatch type and Header	8
6.2. Fragment acknowledgment Dispatch type and Header	8
7. Fragments Recovery	10
8. Forwarding Fragments	11
8.1. Upon the first fragment	12
8.2. Upon the next fragments	13
8.3. Upon the fragment acknowledgments	13
9. Security Considerations	14
10. IANA Considerations	14
11. Acknowledgments	14
12. References	14
12.1. Normative References	14
12.2. Informative References	15
Authors' Addresses	16

1. Introduction

In most Low Power and Lossy Network (LLN) applications, the bulk of the traffic consists of small chunks of data (in the order few bytes to a few tens of bytes) at a time. Given that an 802.15.4 frame can carry 74 bytes or more in all cases, fragmentation is usually not required. However, and though this happens only occasionally, a number of mission critical applications do require the capability to transfer larger chunks of data, for instance to support a firmware upgrades of the LLN nodes or an extraction of logs from LLN nodes. In the former case, the large chunk of data is transferred to the LLN node, whereas in the latter, the large chunk flows away from the LLN node. In both cases, the size can be on the order of 10K bytes or more and an end-to-end reliable transport is required.

Mechanisms such as TCP or application-layer segmentation will be used to support end-to-end reliable transport. One option to support bulk

data transfer over a frame-size-constrained LLN is to set the Maximum Segment Size to fit within the link maximum frame size. Doing so, however, can add significant header overhead to each 802.15.4 frame. This causes the end-to-end transport to be intimately aware of the delivery properties of the underlying LLN, which is a layer violation.

An alternative mechanism combines the use of 6LoWPAN fragmentation in addition to transport or application-layer segmentation. Increasing the Maximum Segment Size reduces header overhead by the end-to-end transport protocol. It also encourages the transport protocol to reduce the number of outstanding datagrams, ideally to a single datagram, thus reducing the need to support out-of-order delivery common to LLNs.

[RFC4944] defines a datagram fragmentation mechanism for LLNs. However, because [RFC4944] does not define a mechanism for recovering fragments that are lost, datagram forwarding fails if even one fragment is not delivered properly to the next IP hop. End-to-end transport mechanisms will require retransmission of all fragments, wasting resources in an already resource-constrained network.

Past experience with fragmentation has shown that missassociated or lost fragments can lead to poor network behavior and, eventually, trouble at application layer. The reader is encouraged to read [RFC4963] and follow the references for more information. That experience led to the definition of the Path MTU discovery [RFC1191] protocol that limits fragmentation over the Internet.

For one-hop communications, a number of media propose a local acknowledgment mechanism that is enough to protect the fragments. In a multihop environment, an end-to-end fragment recovery mechanism might be a good complement to a hop-by-hop MAC level recovery. This draft introduces a simple protocol to recover individual fragments between 6LoWPAN endpoints. Specifically in the case of UDP, valuable additional information can be found in UDP Usage Guidelines for Application Designers [RFC5405].

2. 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 "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and

Goals" [RFC4919] and "Transmission of IPv6 Packets over IEEE 802.15.4 Networks" [RFC4944].

ERP

Error Recovery Procedure.

6LoWPAN endpoints

The LLN nodes in charge of generating or expanding a 6LoWPAN header from/to a full IPv6 packet. The 6LoWPAN endpoints are the points where fragmentation and reassembly take place.

3. Rationale

There are a number of uses for large packets in Wireless Sensor Networks. Such usages may not be the most typical or represent the largest amount of traffic over the LLN; however, the associated functionality can be critical enough to justify extra care for ensuring effective transport of large packets across the LLN.

The list of those usages includes:

Towards the LLN node:

Packages of Commands: A number of commands or a full configuration can be packaged as a single message to ensure consistency and enable atomic execution or complete roll back. Until such commands are fully received and interpreted, the intended operation will not take effect.

Firmware update: For example, a new version of the LLN node software is downloaded from a system manager over unicast or multicast services. Such a reflashing operation typically involves updating a large number of similar LLN nodes over a relatively short period of time.

From the LLN node:

Waveform captures: A number of consecutive samples are measured at a high rate for a short time and then transferred from a sensor to a gateway or an edge server as a single large report.

Data logs: LLN nodes may generate large logs of sampled data for later extraction. LLN nodes may also generate system logs to assist in diagnosing problems on the node or network.

Large data packets: Rich data types might require more than one fragment.

Uncontrolled firmware download or waveform upload can easily result in a massive increase of the traffic and saturate the network.

When a fragment is lost in transmission, all fragments are resent, further contributing to the congestion that caused the initial loss, and potentially leading to congestion collapse.

This saturation may lead to excessive radio interference, or random early discard (leaky bucket) in relaying nodes. Additional queuing and memory congestion may result while waiting for a low power next hop to emerge from its sleeping state.

To demonstrate the severity of the problem, consider a fairly reliable 802.15.4 frame delivery rate of 99.9% over a single 802.15.4 hop. The expected delivery rate of a 5-fragment datagram would be about 99.5% over a single 802.15.4 hop. However, the expected delivery rate would drop to 95.1% over 10 hops, a reasonable network diameter for LLN applications. The expected delivery rate for a 1280-byte datagram is 98.4% over a single hop and 85.2% over 10 hops.

Considering that [RFC4944] defines an MTU is 1280 bytes and that in most incarnations (but 802.15.4G) a 802.15.4 frame can limit the MAC payload to as few as 74 bytes, a packet might be fragmented into at least 18 fragments at the 6LoWPAN shim layer. Taking into account the worst-case header overhead for 6LoWPAN Fragmentation and Mesh Addressing headers will increase the number of required fragments to around 32. This level of fragmentation is much higher than that traditionally experienced over the Internet with IPv4 fragments. At the same time, the use of radios increases the probability of transmission loss and Mesh-Under techniques compound that risk over multiple hops.

4. Requirements

This paper proposes a method to recover individual fragments between LLN endpoints. The method is designed to fit the following requirements of a LLN (with or without a Mesh-Under routing protocol):

Number of fragments

The recovery mechanism must support highly fragmented packets, with a maximum of 32 fragments per packet.

Minimum acknowledgment overhead

Because the radio is half duplex, and because of silent time spent in the various medium access mechanisms, an acknowledgment consumes roughly as many resources as data fragment.

The recovery mechanism should be able to acknowledge multiple fragments in a single message and not require an acknowledgment at all if fragments are already protected at a lower layer.

Controlled latency

The recovery mechanism must succeed or give up within the time boundary imposed by the recovery process of the Upper Layer Protocols.

Support for out-of-order fragment delivery

A Mesh-Under load balancing mechanism such as the ISA100 Data Link Layer can introduce out-of-sequence packets.

The recovery mechanism must account for packets that appear lost but are actually only delayed over a different path.

Optional congestion control

The aggregation of multiple concurrent flows may lead to the saturation of the radio network and congestion collapse.

The recovery mechanism should provide means for controlling the number of fragments in transit over the LLN.

5. Overview

Considering that a multi-hop LLN can be a very sensitive environment due to the limited queuing capabilities of a large population of its nodes, this draft recommends a simple and conservative approach to congestion control, based on TCP congestion avoidance.

Congestion on the forward path is assumed in case of packet loss, and packet loss is assumed upon time out. The draft allows to control the number of outstanding fragments, that have been transmitted but for which an acknowledgment was not received yet. It must be noted that the number of outstanding fragments should not exceed the number of hops in the network, but the way to figure the number of hops is out of scope for this document.

Congestion on the forward path can also be indicated by an Explicit Congestion Notification (ECN) mechanism. Though whether and how ECN [RFC3168] is carried out over the LoWPAN is out of scope, this draft

provides a way for the destination endpoint to echo an ECN indication back to the source endpoint in an acknowledgment message as represented in Figure 5 in Section 6.2.

It must be noted that congestion and collision are different topics. In particular, when a mesh operates on a same channel over multiple hops, then the forwarding of a fragment over a certain hop may collide with the forwarding of a next fragment that is following over a previous hop but in a same interference domain. This draft enables an end-to-end flow control, but leaves it to the sender stack to pace individual fragments within a transmit window, so that a given fragment is sent only when the previous fragment has had a chance to progress beyond the interference domain of this hop. In the case of 6TiSCH [I-D.ietf-6tisch-architecture], which operates over the TimeSlotted Channel Hopping [I-D.ietf-6tisch-tsch] (TSCH) mode of operation of IEEE802.14.5, a fragment is forwarded over a different channel at a different time and it make full sense to fire a next fragment as soon as the previous fragment has had its chance to be forwarded at the next hop, retry (ARQ) operations included.

From the standpoint of a source 6LoWPAN endpoint, an outstanding fragment is a fragment that was sent but for which no explicit acknowledgment was received yet. This means that the fragment might be on the way, received but not yet acknowledged, or the acknowledgment might be on the way back. It is also possible that either the fragment or the acknowledgment was lost on the way.

Because a meshed LLN might deliver frames out of order, it is virtually impossible to differentiate these situations. In other words, from the sender standpoint, all outstanding fragments might still be in the network and contribute to its congestion. There is an assumption, though, that after a certain amount of time, a frame is either received or lost, so it is not causing congestion anymore. This amount of time can be estimated based on the round trip delay between the 6LoWPAN endpoints. The method detailed in [RFC6298] is recommended for that computation.

The reader is encouraged to read through "Congestion Control Principles" [RFC2914]. Additionally [RFC2309] and [RFC5681] provide deeper information on why this mechanism is needed and how TCP handles Congestion Control. Basically, the goal here is to manage the amount of fragments present in the network; this is achieved by to reducing the number of outstanding fragments over a congested path by throttling the sources.

Section 7 describes how the sender decides how many fragments are (re)sent before an acknowledgment is required, and how the sender adapts that number to the network conditions.

6. New Dispatch types and headers

This specification extends "Transmission of IPv6 Packets over IEEE 802.15.4 Networks" [RFC4944] with 4 new dispatch types, for Recoverable Fragments (RFRAG) headers with or without Acknowledgment Request, and for the Acknowledgment back, with or without ECN Echo.

Pattern		Header Type
11	101000	RFRAG - Recoverable Fragment
11	101001	RFRAG-AR - RFRAG with Ack Request
11	101010	RFRAG-ACK - RFRAG Acknowledgment
11	101011	RFRAG-AEC - RFRAG Ack with ECN Echo

Figure 1: Additional Dispatch Value Bit Patterns

In the following sections, the semantics of "datagram_tag", "datagram_offset" and "datagram_size" and the reassembly process are changed from [RFC4944] Section 5.3. "Fragmentation Type and Header." The size and offset are expressed on the compressed packet per [RFC6282] as opposed to the uncompressed - native packet - form.

6.1. Recoverable Fragment Dispatch type and Header




Figure 2: Recoverable Fragment Dispatch type and Header

X: 1 bit; When set, the sender requires an Acknowledgment from the receiver

Sequence: 5 bits; The sequence number of the fragment. Fragments are numbered [0..N] where N is in [0..31].

6.2. Fragment acknowledgment Dispatch type and Header

The specification also defines a 4-octet acknowledgment bitmap that is used to carry selective acknowledgments for the received fragments. A given offset in the bitmap maps one to one with a given sequence number.

The offset of the bit in the bitmap indicates which fragment is acknowledged as follows:




Figure 3: Acknowledgment bitmap encoding

So in the example below Figure 4 it appears that all fragments from sequence 0 to 20 were received but for sequence 1, 2 and 16 that were either lost or are still in the network over a slower path.




Figure 4: Expanding 3 octets encoding

The acknowledgment bitmap is carried in a Fragment Acknowledgment as follows:




Figure 5: Fragment Acknowledgment Dispatch type and Header

Y: 1 bit; Explicit Congestion Notification (ECN) signalling

When set, the sender indicates that at least one of the acknowledged fragments was received with an Explicit Congestion Notification, indicating that the path followed by the fragments is subject to congestion.

acknowledgment Bitmap

An acknowledgment bitmap, whereby but at offset x indicates that fragment x was received.

7. Fragments Recovery

The Recoverable Fragments header RFRAG and RFRAG-AR deprecate the original fragment headers from [RFC4944] and replace them in the fragmented packets. The Fragment Acknowledgment RFRAG-ACK is introduced as a standalone header in message that is sent back to the fragment source endpoint as known by its MAC address. This assumes that the source MAC address in the fragment (is any) and datagram_tag are enough information to send the Fragment Acknowledgment back to the source fragmentation endpoint.

The 6LoWPAN endpoint that fragments the packets at 6LoWPAN level (the sender) controls the Fragment Acknowledgments. It may do that at any fragment to implement its own policy or perform congestion control which is out of scope for this document. When the sender of the fragment knows that an underlying mechanism protects the Fragments already it MAY refrain from using the Acknowledgment mechanism, and never set the Ack Requested bit. The 6LoWPAN endpoint that recomposes the packets at 6LoWPAN level (the receiver) MUST acknowledge the fragments it has received when asked to, and MAY slightly defer that acknowledgment.

The sender transfers a controlled number of fragments and MAY flag the last fragment of a series with an acknowledgment request. The receiver MUST acknowledge a fragment with the acknowledgment request bit set. If any fragment immediately preceding an acknowledgment request is still missing, the receiver MAY intentionally delay its acknowledgment to allow in-transit fragments to arrive. Delaying the acknowledgment might defeat the round trip delay computation so it should be configurable and not enabled by default.

The receiver interacts with the sender using an Acknowledgment message with a bitmap that indicates which fragments were actually received. The bitmap is a 32bit SWORD, which accommodates up to 32 fragments and is sufficient for the 6LoWPAN MTU. For all n in [0..31], bit n is set to 1 in the bitmap to indicate that fragment with sequence n was received, otherwise the bit is set to 0. All zeros is a NULL bitmap that indicates that the fragmentation process was canceled by the receiver for that datagram.

The receiver MAY issue unsolicited acknowledgments. An unsolicited acknowledgment enables the sender endpoint to resume sending if it had reached its maximum number of outstanding fragments or indicate that the receiver has cancelled the process of an individual datagram. Note that acknowledgments might consume precious resources

so the use of unsolicited acknowledgments should be configurable and not enabled by default.

The sender arms a retry timer to cover the fragment that carries the Acknowledgment request. Upon time out, the sender assumes that all the fragments on the way are received or lost. The process must have completed within an acceptable time that is within the boundaries of upper layer retries. The method detailed in [RFC6298] is recommended for the computation of the retry timer. It is expected that the upper layer retries obey the same or friendly rules in which case a single round of fragment recovery should fit within the upper layer recovery timers.

Fragments are sent in a round robin fashion: the sender sends all the fragments for a first time before it retries any lost fragment; lost fragments are retried in sequence, oldest first. This mechanism enables the receiver to acknowledge fragments that were delayed in the network before they are actually retried.

When the sender decides that a packet should be dropped and the fragmentation process canceled, it sends a pseudo fragment with the datagram_offset, sequence and datagram_size all set to zero, and no data. Upon reception of this message, the receiver should clean up all resources for the packet associated to the datagram_tag. If an acknowledgment is requested, the receiver responds with a NULL bitmap.

The receiver might need to cancel the process of a fragmented packet for internal reasons, for instance if it is out of recomposition buffers, or considers that this packet is already fully recomposed and passed to the upper layer. In that case, the receiver SHOULD indicate so to the sender with a NULL bitmap. Upon an acknowledgment with a NULL bitmap, the sender MUST drop the datagram.

8. Forwarding Fragments

This specification enables intermediate routers to forward fragments with no intermediate reconstruction of the entire packet. Upon the first fragment, the routers lay an label along the path that is followed by that fragment (that is IP routed), and all further fragments are label switched along that path. As a consequence, alternate routes not possible for individual fragments. The datagram_tag is used to carry the label, that is swapped at each hop.

8.1. Upon the first fragment

In route over the L2 source changes at each hop. The label that is formed and placed in the datagram_tag is associated to the source MAC and only valid (and unique) for that source MAC. Say the first fragment has:

Source IPv6 address = IP_A (maybe hops away)

Destination IPv6 address = IP_B (maybe hops away)

Source MAC = MAC_prv (prv as previous)

Datagram_tag= DT_prv

The intermediate router that forwards individual fragments does the following:

a route lookup to get Next hop IPv6 towards IP_B, which resolves as IP_nxt (nxt as next)

a MAC address resolution to get the MAC address associated to IP_nxt, which resolves as MAC_nxt

Since it is a first fragment of a packet from that source MAC address MAC_prv for that tag DT_prv, the router:

cleans up any leftover resource associated to the tuple (MAC_prv, DT_prv)

allocates a new label for that flow, DT_nxt, from a Least Recently Used pool or some similar procedure.

allocates a Label swap structure indexed by (MAC_prv, DT_prv) that contains (MAC_nxt, DT_nxt)

allocates a Label swap structure indexed by (MAC_nxt, DT_nxt) that contains (MAC_prv, DT_prv)

swaps the MAC info to from self to MAC_nxt

Swaps the datagram_tag to DT_nxt

At this point the router is all set and can forward the packet to nxt.

8.2. Upon the next fragments

Upon next fragments (that are not first fragment), the router expects to have already Label swap structure indexed by (MAC_prv, DT_prv). The router:

lookups up the Label swap entry for (MAC_prv, DT_prv), which resolves as (MAC_nxt, DT_nxt)

swaps the MAC info to from self to MAC_nxt;

Swaps the datagram_tag to DT_nxt

At this point the router is all set and can forward the packet to nxt.

if the Label swap entry for (MAC_src, DT_src) is not found, the router builds an RFRAG-ACK to indicate the error. The acknowledgment message has the following information:

MAC info set to from self to MAC_prv as found in the fragment

Swaps the datagram_tag set to DT_prv

Bitmap of all zeroes to indicate the error

At this point the router is all set and can send the RFRAG-ACK back to the previous router.

8.3. Upon the fragment acknowledgments

Upon fragment acknowledgments next fragments (that are not first fragment), the router expects to have already Label swap structure indexed by (MAC_nxt, DT_nxt). The router:

lookups up the Label swap entry for (MAC_nxt, DT_nxt), which resolves as (MAC_prv, DT_prv)

swaps the MAC info to from self to MAC_prv;

Swaps the datagram_tag to DT_prv

At this point the router is all set and can forward the RFRAG-ACK to prv.

if the Label swap entry for (MAC_nxt, DT_nxt) is not found, it simply drops the packet.

if the RFRAG-ACK indicates either an error or that the fragment was fully receive, the router schedules the Label swap entries for recycling. If the RFRAG-ACK is lost on the way back, the source may retry the last fragments, which will result as an error RFRAG-ACK from the first router on the way that has already cleaned up.

9. Security Considerations

The process of recovering fragments does not appear to create any opening for new threat compared to "Transmission of IPv6 Packets over IEEE 802.15.4 Networks" [RFC4944].

10. IANA Considerations

Need extensions for formats defined in "Transmission of IPv6 Packets over IEEE 802.15.4 Networks" [RFC4944].

11. Acknowledgments

The author wishes to thank Jay Werb, Christos Polyzois, Soumitri Kolavennu, Pat Kinney, Margaret Wasserman, Richard Kelsey, Carsten Bormann and Harry Courtice for their contributions and review.

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>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<http://www.rfc-editor.org/info/rfc4944>>.
- [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>>.
- [RFC6298] Paxson, V., Allman, M., Chu, J., and M. Sargent, "Computing TCP's Retransmission Timer", RFC 6298, DOI 10.17487/RFC6298, June 2011, <<http://www.rfc-editor.org/info/rfc6298>>.

12.2. Informative References

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

Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", draft-ietf-6tisch-architecture-10 (work in progress), June 2016.

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

Watteyne, T., Palattella, M., and L. Grieco, "Using IEEE802.15.4e TSCH in an IoT context: Overview, Problem Statement and Goals", draft-ietf-6tisch-tsch-06 (work in progress), March 2015.

[RFC1191] Mogul, J. and S. Deering, "Path MTU discovery", RFC 1191, DOI 10.17487/RFC1191, November 1990, <<http://www.rfc-editor.org/info/rfc1191>>.

[RFC2309] Braden, B., Clark, D., Crowcroft, J., Davie, B., Deering, S., Estrin, D., Floyd, S., Jacobson, V., Minshall, G., Partridge, C., Peterson, L., Ramakrishnan, K., Shenker, S., Wroclawski, J., and L. Zhang, "Recommendations on Queue Management and Congestion Avoidance in the Internet", RFC 2309, DOI 10.17487/RFC2309, April 1998, <<http://www.rfc-editor.org/info/rfc2309>>.

[RFC2914] Floyd, S., "Congestion Control Principles", BCP 41, RFC 2914, DOI 10.17487/RFC2914, September 2000, <<http://www.rfc-editor.org/info/rfc2914>>.

[RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, DOI 10.17487/RFC3168, September 2001, <<http://www.rfc-editor.org/info/rfc3168>>.

[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>>.

[RFC4963] Heffner, J., Mathis, M., and B. Chandler, "IPv4 Reassembly Errors at High Data Rates", RFC 4963, DOI 10.17487/RFC4963, July 2007, <<http://www.rfc-editor.org/info/rfc4963>>.

- [RFC5405] Eggert, L. and G. Fairhurst, "Unicast UDP Usage Guidelines for Application Designers", BCP 145, RFC 5405, DOI 10.17487/RFC5405, November 2008, <<http://www.rfc-editor.org/info/rfc5405>>.
- [RFC5681] Allman, M., Paxson, V., and E. Blanton, "TCP Congestion Control", RFC 5681, DOI 10.17487/RFC5681, September 2009, <<http://www.rfc-editor.org/info/rfc5681>>.

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

Jonathan W. Hui
Nest Labs
3400 Hillview Ave
Palo Alto, California 94304
USA

Email: jonhui@nestlabs.com