

6lo
Internet-Draft
Updates: 8505 (if approved)
Intended status: Standards Track
Expires: August 29, 2019

P. Thubert, Ed.
Cisco
B. Sarikaya

M. Sethi
Ericsson
R. Struik
Struik Security Consultancy
February 25, 2019

Address Protected Neighbor Discovery for Low-power and Lossy Networks
draft-ietf-6lo-ap-nd-11

Abstract

This document specifies an extension to 6LoWPAN Neighbor Discovery (ND) protocol defined in RFC6775 and updated in RFC8505. 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 August 29, 2019.

Copyright Notice

Copyright (c) 2019 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. References	4
2.3. Abbreviations	5
3. Updating RFC 8505	6
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. Nonce Option	9
4.5. NDP Signature Option	9
5. Protocol Scope	10
6. Protocol Flows	11
6.1. First Exchange with a 6LR	11
6.2. NDPSO generation and verification	13
6.3. Multihop Operation	14
7. Security Considerations	16
7.1. Inheriting from RFC 3971	16
7.2. Related to 6LoWPAN ND	17
7.3. ROVR Collisions	17
7.4. Implementation Attacks	17
7.5. Cross-Protocol Attacks	18
8. IANA considerations	18
8.1. CGA Message Type	18
8.2. Crypto-Type Subregistry	18
9. Acknowledgments	19
10. References	19
10.1. Normative References	19
10.2. Informative references	20
Appendix A. Requirements Addressed in this Document	22

Appendix B. Representation Conventions	22
B.1. Signature Schemes	22
B.2. Integer Representation for ECDSA signatures	23
B.3. Alternative Representations of Curve25519	23
Authors' Addresses	25

1. Introduction

Neighbor Discovery Optimizations for 6LoWPAN networks [RFC6775] (6LoWPAN ND) adapts the original IPv6 neighbor discovery (NDv6) 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. 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 6LoWPAN ND [RFC6775] 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). ROVR is defined in [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 the an attacker to steal the address and redirect traffic for that address. [RFC8505] defines an Extended Address Registration Option (EARO) option that allows to transport 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 enables Source Address Validation (SAVI) [RFC7039]. This ensures

that only the actual owner uses a registered address in the IPv6 source address field. A 6LN can only use a 6LR for forwarding packets only if it has previously registered the address used in the source field of the IPv6 packet.

The 6lo adaptation layer in [RFC4944] and [RFC6282] requires a device to form its IPv6 addresses based on its Layer-2 address to enable a better compression. This is incompatible with Secure Neighbor Discovery (SeND) [RFC3971] and Cryptographically Generated Addresses (CGAs) [RFC3972], since they derive the Interface ID (IID) in IPv6 addresses with cryptographic keys.

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. References

Terms and concepts from the following documents are used in this specification:

- o "Terms Used in Routing for Low-Power and Lossy Networks (LLNs)" [RFC7102],
- o "SEcure Neighbor Discovery (SEND)" [RFC3971],
- o "Cryptographically Generated Addresses (CGA)" [RFC3972],
- o "Neighbor Discovery for IP version 6" [RFC4861] ,
- o "IPv6 Stateless Address Autoconfiguration" [RFC4862],
- o "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals " [RFC4919],
- o "Neighbor Discovery Optimization for Low-power and Lossy Networks" [RFC6775], and
- o "Registration Extensions for 6LoWPAN Neighbor Discovery" [RFC8505].

2.3. Abbreviations

This document uses the following abbreviations:

6BBR: 6LoWPAN Backbone Router
6LBR: 6LoWPAN Border Router
6LN: 6LoWPAN Node
6LR: 6LoWPAN Router
6CIO: Capability Indication Option
ARO: Address Registration Option
CIPO: Crypto-ID Parameters Option
LLN: Low-Power and Lossy Network
NA: Neighbor Advertisement
NCE: Neighbor Cache Entry
ND: Neighbor Discovery
NDP: Neighbor Discovery Protocol
NDPSO: NDP Signature Option
NS: Neighbor Solicitation
ROVR: Registration Ownership Verifier
RPL: IPv6 Routing Protocol for LLNs
RA: Router Advertisement
RS: Router Solicitation
RSAO: RSA Signature Option
TID: Transaction ID

3. Updating RFC 8505

This specification introduces a new token called a cryptographic identifier (Crypto-ID) that is used to prove indirectly the ownership of an address that is being registered by means of [RFC8505].

In order to prove its ownership of a Crypto-ID, the registering node needs to supply certain parameters including a nonce and a signature that will prove that the node has the private-key corresponding to the public-key used to build the Crypto-ID. This specification adds the capability to carry new options in the NS(EARO) and the NA(EARO). The NS(EARO) carries a variation of the CGA Option (Section 4.3), a Nonce option and a variation of the RSA Signature option (Section 4.5) in the NS(EARO). The NA(EARO) carries a Nonce option.

4. New Fields and Options

In order to avoid the need for new ND option types, this specification reuses and extends options defined in SEND [RFC3971] and 6LoWPAN ND [RFC6775] [RFC8505]. This applies in particular to the CGA option and the RSA Signature Option. This specification provides aliases for the specific variations of those options as used in this document. The presence of the EARO option in the NS/NA messages indicates that the options are to be processed as specified in this document, and not as defined in SEND [RFC3971].

4.1. New Crypto-ID

The Crypto-ID can be used as a replacement to the MAC address in the ROVR field of the EARO option and the EDAR message, and is associated with the Registered Address. 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 computation of the Crypto-ID requires the support of Elliptic Curve Cryptography (ECC) and that of a hash function as detailed in Section 6.2. The elliptic curves and the hash functions that can be used with this specification are listed in Table 1 in Section 8.2. The signature scheme that specifies which combination is used is signaled by a Crypto-Type in a new Crypto-ID Parameters Option (see Section 4.3).

4.2. Updated EARO

This specification updates the EARO option as follows:




Figure 1: Enhanced Address Registration Option

- Type: 33
- Length: 8-bit unsigned integer. The length of the option (including the type and length fields) in units of 8 bytes.
- Status: 8-bit unsigned integer. Indicates the status of a registration in the NA response. MUST be set to 0 in NS messages.
- Opaque: Defined in [RFC8505].
- Rsvd (Reserved): This field is unused. It MUST be initialized 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, and 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]. No other new Status values are defined.

4.3. Crypto-ID Parameters Option

This specification defines the Crypto-ID Parameters Option (CIPO), as a variation of the CGA Option that carries the parameters used to form a Crypto-ID. In order to provide cryptographic agility [RFC7696], this specification supports different elliptic curves, indicated by a Crypto-Type field. NIST P-256 [FIPS186-4] MUST be supported by all implementations. The Edwards-Curve Digital Signature Algorithm (EdDSA) curve Ed25519 (PureEdDSA) [RFC8032] MAY be supported as an alternate.

The type of cryptographic algorithm used in the calculation of the Crypto-ID is signaled by the Crypto-Type field of the CIPO as specified in Table 1 in Section 8.2. Although the different signature schemes target similar cryptographic strength, they rely on different curves, hash functions, signature algorithms, and/or representation conventions.




Figure 2: Crypto-ID Parameters Option

Type: 11. This is the same value as the CGA Option, CIPO is a particular case of the CGA option

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

Modifier: 8-bit unsigned integer.

Pad Length: 8-bit unsigned integer. The length of the Padding field.

Crypto-Type: The type of cryptographic algorithm used in calculation Crypto-ID (see Table 1 in Section 8.2).

Public Key: JWK-Encoded Public Key [RFC7517].

Padding: A variable-length field making the option length a multiple of 8, containing as many octets as specified in the Pad Length field.

The implementation of multiple hash functions in a constrained devices may consume excessive amounts of program memory. [I-D.ietf-lwig-curve-representations] 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. Nonce Option

This document reuses the Nonce Option defined in section 5.3.2. of SEND [RFC3971] without a change.

4.5. NDP Signature Option

This document reuses the RSA Signature Option (RSAO) defined in section 5.2. of SEND [RFC3971]. Admittedly, the name is ill-chosen since the option is extended for non-RSA Signatures and this specification defines an alias to avoid the confusion.

The description of the operation on the option detailed in section 5.2. of SEND [RFC3971] apply, but for the following changes:

- o The 128-bit CGA Message Type tag [RFC3972] for AP-ND is 0x8701 55c8 0cca dd32 6ab7 e415 f148 84d0. (The tag value has been generated by the editor of this specification on random.org).
- o The signature is computed using the hash algorithm and the digital signature indicated in the Crypto-Type field of the CIP0 option

using the private-key corresponding the public-key passed in the CIP0.

- o The alias NDP Signature Option (NDPSO) can be used to refer to the RSAO when used as described in this specification.

5. Protocol Scope

The scope of the protocol specified here is a 6LoWPAN Low Power Lossy Network (LLN), typically a stub network connected to a larger IP network via a Border Router called a 6LBR per [RFC6775]. A 6LBR has sufficient capability to satisfy the needs of duplicate address detection.

The 6LBR maintains registration state for all devices in its attached LLN. Together with the first-hop router (the 6LR), the 6LBR assures uniqueness and grants ownership of an IPv6 address before it can be used in the LLN. This is in contrast to a traditional network that relies on IPv6 address auto-configuration [RFC4862], where there is no guarantee of ownership from the network, and each IPv6 Neighbor Discovery packet must be individually secured [RFC3971].




Figure 3: Basic Configuration

In a mesh network, the 6LR is directly connected to the host device. This specification mandates that the peer-wise layer-2 security is deployed so that all the packets from a particular host are securely identifiable by the 6LR. The 6LR may be multiple hops away from the 6LBR. Packets are routed between the 6LR and the 6LBR via other 6LRs. This specification mandates that a chain of trust is established so that a packet that was validated by the first 6LR can be safely routed by other on-path 6LRs to the 6LBR.

6. Protocol Flows

The 6LR/6LBR ensures first-come/first-serve by storing the EARO information including the Crypto-ID associated to the node being registered. The node can claim any address as long as it is the first to make such a claim. After a successful registration, the node becomes the owner of the registered address and the address is bound to the Crypto-ID in the 6LR/6LBR registry.

This specification enables the 6LR to verify the ownership of the binding at any time assuming that the "C" flag is set. The verification prevents other nodes from stealing the address and trying to attract traffic for that address or use it as their source address.

A node may use multiple IPv6 addresses at the same time. The node may use a same Crypto-ID, to prove the ownership of multiple IPv6 addresses. The separation of the address and the cryptographic material avoids the constrained device to compute multiple keys for multiple addresses. The registration process allows the node to use the same Crypto-ID for all of its addresses.

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. The on-link (local) protocol interactions are shown in Figure 4. 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 4). The Nonce option MUST contain a random Nonce value that was never used with this device.

The 6LN replies to the challenge with an NS(EARO) that includes a new Nonce option (shown as NonceLN in Figure 4), the CIPO (Section 4.3), and the NDPSO containing the signature. The information associated to a Crypto-ID stored by the 6LR on the first NS exchange where it appears. The 6LR MUST store the CIPO parameters associated with the Crypto-ID so it can be used for more than one address.

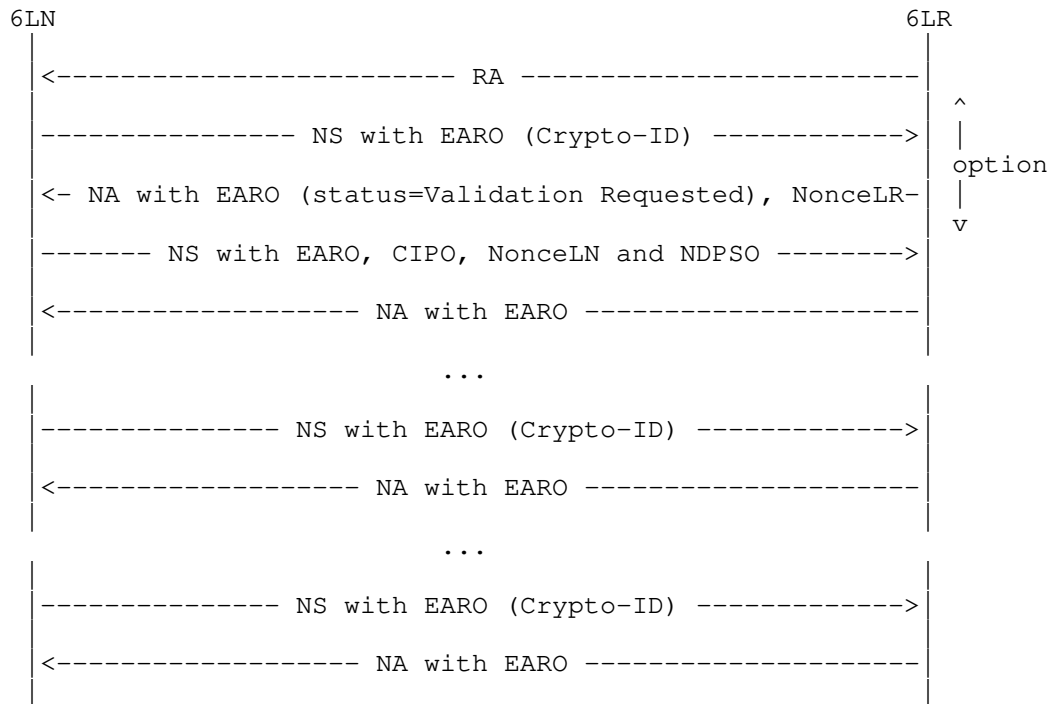


Figure 4: On-link Protocol Operation

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

- o 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, it SHOULD challenge by responding with a NA(EARO) with a status of "Validation Requested".
- o The challenge is triggered when the registration for a Source Link-Layer Address is not verifiable either at the 6LR or the 6LBR. In the latter case, the 6LBR returns a status of "Validation Requested" in the DAR/DAC exchange, which is echoed by the 6LR in the NA (EARO) back to the registering node. The challenge MUST NOT alter a valid registration in the 6LR or the 6LBR.
- o Upon receiving a NA(EARO) with a status of "Validation Requested", 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 4.5) option that proves its ownership of the Crypto-ID and intent of registering the Target Address.

- o 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. It also verifies the signature contained in the NDPSO option. If the Crypto-ID does not match with the public-key in the CIPO option, or if the signature in the NDPSO option cannot be verified, the validation fails.
- o 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 to register an alternate target 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:

- o Concatenate the following in the order listed:
 1. 128-bit type tag (in network byte order)
 2. JWK-encoded public key
 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 random nonce is at least 6 bytes long as defined in [RFC3971].
 5. NonceLN sent from the 6LN (in network byte order). The random nonce is at least 6 bytes long as defined in [RFC3971].
 6. The length of the ROVR field in the NS message containing the Crypto-ID that was sent.
 7. 1-byte (in network byte order) Crypto-Type value sent in the CIPO option.
- o Depending on the Crypto-Type, apply the hash function on this concatenation.

- o Depending on the Crypto-Type, sign the hash output with ECDSA (if curve P-256 is used) or sign the hash with EdDSA (if curve Ed25519 (PureEdDSA)).

The 6LR on receiving the NDPSO and CIPO options first hashes the JWK encoded public-key in the CIPO option to make sure that the leftmost bits up to the size of the ROVR match. Only if the check is successful, it tries to verify the signature in the NDPSO option using the following.

- o Concatenate the following in the order listed:
 1. 128-bit type tag (in network byte order)
 2. JWK-encoded public key received in the CIPO option
 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 random 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 random nonce is at least 6 bytes long as defined in [RFC3971].
 6. The length of the ROVR field in the NS message containing the Crypto-ID that was received.
 7. 1-byte (in network byte order) Crypto-Type value received in the CIPO option.
- o Depending on the Crypto-Type indicated by the (6LN) in the CIPO, apply the hash function on this concatenation.
- o Verify the signature with the public-key received and the locally computed values. If the verification succeeds, the 6LR and 6LBR add the state information about the Crypto-ID, public-key and Target Address being registered to their database.

6.3. Multihop Operation

In a multihop 6LoWPAN, the registration with Crypto-ID is propagated to 6LBR as described in this section. If the 6LR and the 6LBR maintain a security association, then there is no need to propagate the proof of ownership to the 6LBR.

A new device 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]. The 6LR validates the address with an 6LBR using a DAR/DAC exchange, and the 6LR confirms (or denies) the address ownership with an NA message that also carries an Address Registration Option.

Figure 5 illustrates a registration flow all the way to a 6LowPAN Backbone Router (6BBR) [I-D.ietf-6lo-backbone-router].

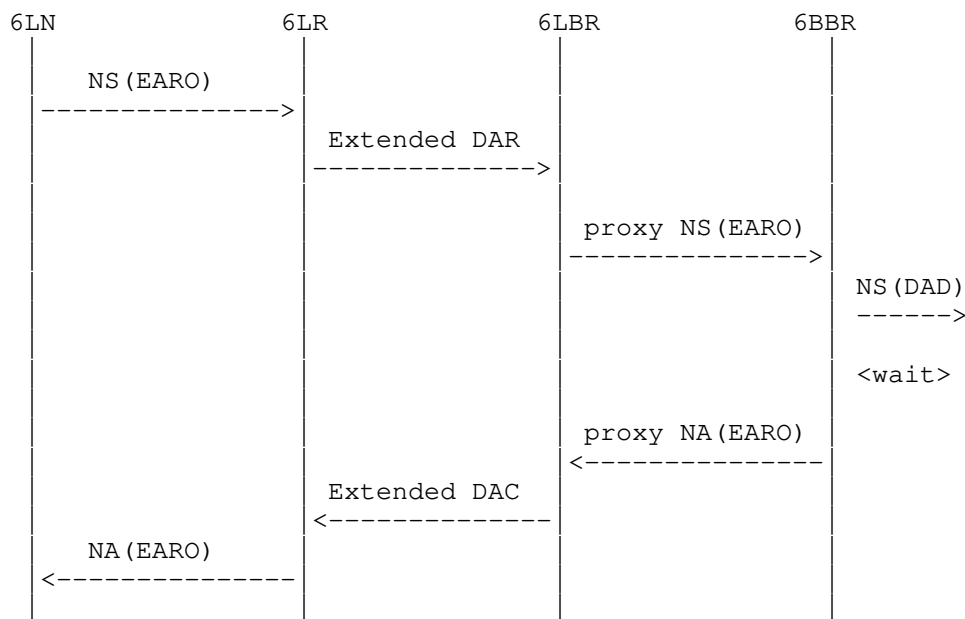


Figure 5: (Re-)Registration Flow

In a multihop 6LoWPAN, a 6LBR sends RAs with prefixes downstream and the 6LR receives and relays them to the nodes. 6LR and 6LBR communicate using ICMPv6 Duplicate Address Request (DAR) and Duplicate Address Confirmation (DAC) messages. The DAR and DAC use the same message format as NS and NA, but have different ICMPv6 type values.

In AP-ND we extend DAR/DAC messages to carry cryptographically generated ROVR. In a multihop 6LoWPAN, the node exchanges the messages shown in Figure 5. The 6LBR must identify who owns an address (EUI-64) to defend it, if there is an attacker on another 6LR.

7. Security Considerations

7.1. 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. 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 and is thus the best candidate to validate the registration for the device attached to it. 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

Nonces (NonceLR and NonceLN) generated by the 6LR and 6LN guarantees against replay attacks of the NS(EARO).

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.

7.2. Related to 6LoWPAN ND

The threats discussed in 6LoWPAN ND [RFC6775][RFC8505] also apply here. Compared with SeND, 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. SeND forces the IPv6 address to be cryptographic since it integrates the CGA as the IID in the IPv6 address. This specification frees the device to form its addresses in any fashion, thereby enabling not only 6LoWPAN compression which derives IPv6 addresses from Layer-2 addresses but also privacy addresses.

7.3. 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. The formula for calculating the probability of a collision is $1 - e^{-k^2/(2n)}$ where n is the maximum population size (2^{64} here, 1.84E19) and K is the actual population (number of nodes). If the Crypto-ID is 64-bits (the least possible size allowed), the chance of a collision is 0.01% when the network contains 66 million nodes. Moreover, the collision is only relevant when this happens within one stub network (6LBR). In the case of such a collision, an attacker may be able to claim the registered address of another legitimate node. However for this to happen, the attacker would also need to know the address which was registered by the legitimate node. This registered address is never broadcasted on the network and therefore providing an additional 64-bits that an attacker must correctly guess. To prevent address disclosure, it is RECOMMENDED that nodes derive the address being registered independently of the ROVR.

7.4. 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 SHA-256 used with ECDSA.

7.5. Cross-Protocol Attacks

The same private key **MUST NOT** be reused with more than one signature scheme in this specification.

8. IANA considerations

8.1. CGA Message Type

This document defines a new 128-bit value 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, and Representation Conventions, as shown in Table 1, which together specify a signature scheme. 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	Edwards, compressed, LSB/lsw first	Weierstrass, (un)compressed, MSB/msb first
Defining specification	RFC THIS	RFC THIS	RFC THIS

Table 1: Crypto-Types

New Crypto-Type values providing similar or better security (with less code) may be defined in the future.

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

9. Acknowledgments

Many thanks to Charlie Perkins for his in-depth review and constructive suggestions. We are also especially grateful to Robert Moskowitz for his comments that led to many improvements.

10. References

10.1. Normative References

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

- [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>>.
- [RFC7517] Jones, M., "JSON Web Key (JWK)", RFC 7517, DOI 10.17487/RFC7517, May 2015, <<https://www.rfc-editor.org/info/rfc7517>>.
- [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>>.
- [SEC1] SEC1, "SEC 1: Elliptic Curve Cryptography, Version 2.0", Standards for Efficient Cryptography , June 2009.

10.2. Informative references

- [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>.
- [I-D.ietf-6lo-backbone-router] Thubert, P., Perkins, C., and E. Levy-Abegnoli, "IPv6 Backbone Router", draft-ietf-6lo-backbone-router-11 (work in progress), February 2019.
- [I-D.ietf-lwig-curve-representations] Struik, R., "Alternative Elliptic Curve Representations", draft-ietf-lwig-curve-representations-01 (work in progress), November 2018.
- [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>>.
- [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>>.
- [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>>.
- [RFC7102] Vasseur, JP., "Terms Used in Routing for Low-Power and Lossy Networks", RFC 7102, DOI 10.17487/RFC7102, January 2014, <<https://www.rfc-editor.org/info/rfc7102>>.
- [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>>.
- [RFC7696] 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>>.
- [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>>.
- [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>>.

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.

- o 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.
- o 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.
- o 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.
- o 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.
- o The Address Registration Option used in the ND registration SHOULD be extended to carry the relevant forms of Unique Interface Identifier.
- o 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 integer 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 (see Appendix B.3) 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 and on the encoding conversions, 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 a corresponding curve in short-Weierstrass form (see Appendix B.3) 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 integer encoding, see Appendix B.2

B.2. Integer Representation for ECDSA signatures

With ECDSA, each signature is a pair (r, s) of integers [FIPS186-4]. Each integer is encoded as a fixed-size 256-bit bit string, where each integer is represented according to the Field Element to Octet String and Octet String to Bit String conversion rules in [SEC1] and where the ordered pair of integers is represented as the rightconcatenation of the resulting representation values. The inverse operation follows the corresponding Bit String to Octet String and Octet String to Field Element conversion rules of [SEC1].

B.3. 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 details of the coordinate transformation referenced above, see [RFC7748] and [I-D.ietf-lwig-curve-representations].

General parameters (for all curve models):

p $2^{255}-19$

(=0x7ffffffff ffffffff ffffffff ffffffff ffffffff ffffffff
ffffffff ffffffff)

h 8

n 72370055773322622139731865630429942408571163593799076060019509382
85454250989

(= 2^{252} + 0x14def9de a2f79cd6 5812631a 5cf5d3ed)

Montgomery curve-specific parameters (for Curve25519):

A 486662

B 1

Gu 9 (=0x9)

Gv 14781619447589544791020593568409986887264606134616475288964881837
755586237401

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

Twisted Edwards curve-specific parameters (for Edwards25519):

a -1 (-0x01)

d -121665/121666

(=370957059346694393431380835087545651895421138798432190163887855
33085940283555)

(=0x52036cee 2b6ffe73 8cc74079 7779e898 00700a4d 4141d8ab
75eb4dca 135978a3)

Gx 15112221349535400772501151409588531511454012693041857206046113283
949847762202

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

Gy 4/5

(=463168356949264781694283940034751631413079938662562256157830336
03165251855960)

(=0x66666666 66666666 66666666 66666666 66666666 66666666
66666666 66666658)

Weierstrass curve-specific parameters (for Wei25519):

a 19298681539552699237261830834781317975544997444273427339909597334
573241639236

(=0x2aaaaaaaa aaaaaaaaa aaaaaaaaa aaaaaaaaa aaaaaaaaa aaaaaaaaa
aaaaaaaa98 4914a144)

b 55751746669818908907645289078257140818241103727901012315294400837
956729358436

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

GX 19298681539552699237261830834781317975544997444273427339909597334
652188435546

(=0x2aaaaaaaa aaaaaaaaa aaaaaaaaa aaaaaaaaa aaaaaaaaa aaaaaaaaa
aaaaaaaa aaad245a)

GY 14781619447589544791020593568409986887264606134616475288964881837
755586237401

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

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

Behcet Sarikaya

Email: sarikaya@ieee.org

Mohit Sethi
Ericsson
Jorvas 02420
Finland

Email: mohit@piuha.net

Rene Struik
Struik Security Consultancy

Email: rstruik.ext@gmail.com

6lo
Internet-Draft
Updates: 6775, 8505 (if approved)
Intended status: Standards Track
Expires: August 8, 2019

P. Thubert, Ed.
Cisco Systems
C. Perkins
Futurewei
E. Levy-Abegnoli
Cisco Systems
February 4, 2019

IPv6 Backbone Router
draft-ietf-6lo-backbone-router-11

Abstract

This document updates RFC 4861 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 MultiLink 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 August 8, 2019.

Copyright Notice

Copyright (c) 2019 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	4
2.1. BCP 14	4
2.2. New Terms	5
2.3. Acronym Definitions	6
2.4. References	7
3. Overview	7
3.1. Updating RFC 6775 and RFC 8505	9
3.2. Access Link	10
3.3. Route-Over Mesh	11
3.4. The Binding Table	12
3.5. Primary and Secondary 6BBRs	13
3.6. Using Optimistic DAD	14
4. MultiLink Subnet Considerations	14
5. Optional 6LBR serving the MultiLink Subnet	15
6. Using IPv6 ND Over the Backbone Link	15
7. Routing Proxy Operations	16
8. Bridging Proxy Operations	17
9. Creating and Maintaining a Binding	18
9.1. Operation on a Binding in Tentative State	19
9.2. Operation on a Binding in Reachable State	20
9.3. Operation on a Binding in Stale State	21
10. Registering Node Considerations	22
11. Security Considerations	23
12. Protocol Constants	23
13. IANA Considerations	23
14. Acknowledgments	23
15. References	24
15.1. Normative References	24
15.2. Informative References	25
Appendix A. Possible Future Extensions	28
Appendix B. Applicability and Requirements Served	28
Authors' Addresses	30

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 node (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.

Like 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, IPv6 addresses very rarely conflict because of the entropy of the 64-bit Interface IDs, not because address duplications are detected and resolved.

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 broadcasted 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 total 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 routes between subnets, or even by assigning a /64 prefix to each wireless node (see [RFC8273]).

Another way is to proxy at the boundary of the wired and wireless domains the Layer-3 protocols that rely on MAC Layer broadcast operations. For instance, IEEE 802.11 [IEEEstd80211] situates proxy-ARP (IPv4) and proxy-ND (IPv6) functions at the Access Points (APs). The 6BBR provides a proxy-ND function and can be extended for proxy-ARP in a continuation specification.

IPv6 proxy-ND services can be obtained by snooping the IPv6 ND protocol (see [I-D.bi-savi-wlan]). Proprietary techniques for IPv6 ND and DHCP snooping have been used; although snooping does eliminate undesirable broadcast transmissions, 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.

This specification defines the 6BBR as a Routing Registrar [RFC8505] that provide proxy services for IPv6 Neighbor Discovery. Backbone Routers federate multiple LLNs over a Backbone Link to form a MultiLink Subnet (MLSN). Backbone Routers placed along the LLN edge of the Backbone handle IPv6 Neighbor Discovery, and forward packets on behalf of registered nodes.

An LLN node (6LN) registers all its IPv6 Addresses using an NS(EOA) as specified in [RFC8505] to the 6BBR. The 6BBR is also a Border Router that 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. Additional benefits are discussed in Appendix B.

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 MultiLink Subnet. The proxy-ND operation of 6BBRs over the Backbone and the access links provides the appearance of a subnet for IPv6 ND.

Sleeping Proxy

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

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 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. The Bridging Proxy advertises the MAC Address of the Registering Node as the TLLA in the proxied NAs over the Backbone. In that case, the MAC Address and the mobility of 6LN is still visible across the bridged Backbone, and the 6BR may be configured to proxy for Link Local Addresses.

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. Acronym Definitions

This document uses the following acronyms:

6BBR: 6LoWPAN Backbone Router

6LBR: 6LoWPAN Border Router

6LN: 6LoWPAN Node

6LR: 6LoWPAN Router

6CIO: Capability Indication Option

ARO: Address Registration Option

DAC: Duplicate Address Confirmation

DAD: Duplicate Address Detection

DAR: Duplicate Address Request

EDAC: Extended Duplicate Address Confirmation

EDAR: Extended Duplicate Address Request

DODAG: Destination-Oriented Directed Acyclic Graph

LLN: Low-Power and Lossy Network

NA: Neighbor Advertisement

NCE: Neighbor Cache Entry

ND: Neighbor Discovery

NDP: Neighbor Discovery Protocol

NS: Neighbor Solicitation

ROVR: Registration Ownership Verifier

RPL: IPv6 Routing Protocol for LLNs

RA: Router Advertisement

RS: Router Solicitation

TID: Transaction ID

2.4. References

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

- o "Neighbor Discovery for IP version 6" [RFC4861], "IPv6 Stateless Address Autoconfiguration" [RFC4862] and "Optimistic Duplicate Address Detection" [RFC4429],
- o "Neighbor Discovery Proxies (proxy-ND)" [RFC4389] and "MultiLink Subnet Issues" [RFC4903],
- o "Problem Statement and Requirements for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing" [RFC6606], and
- o Neighbor Discovery Optimization for Low-Power and Lossy Networks [RFC6775] and "Registration Extensions for 6LoWPAN Neighbor Discovery" [RFC8505].

3. Overview

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

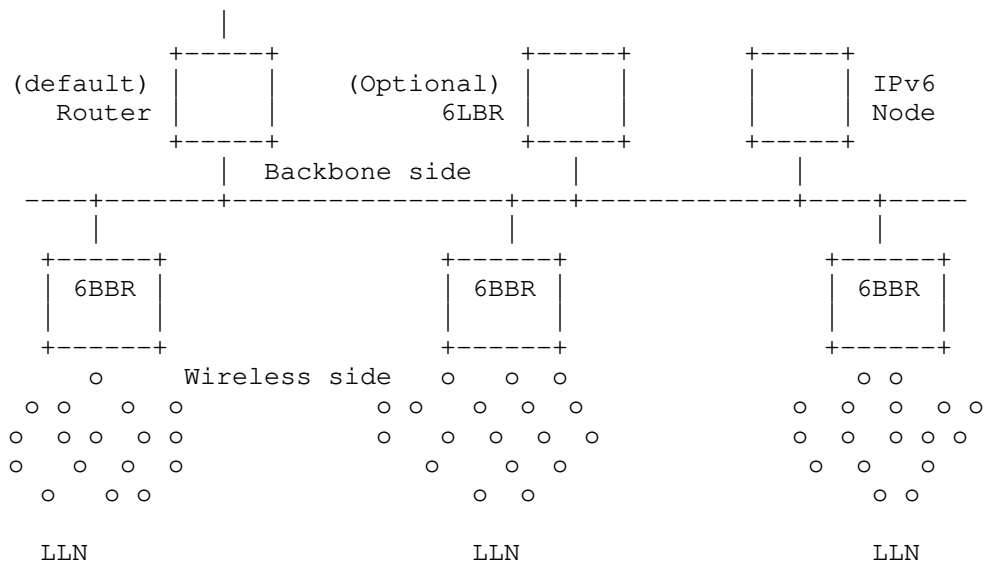


Figure 1: Backbone Link and Backbone Routers

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:

- o Multilink-subnet functions (provided by the 6BBR on the backbone) performed on behalf of registered 6LNs, and
- o 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 Nodes called a Binding Table. The combined Binding Tables of all the 6BBRs on a backbone form a distributed database of 6LNs that reside in the LLNs or on the IPv6 Backbone.

Unless otherwise configured, a 6BBR does the following:

- o Create a new entry in a Binding Table for a new Registered Address and ensure that the Address is not duplicated over the Backbone
- o Defend a Registered Address over the Backbone using NA messages on behalf of the sleeping 6LN
- o Advertise a Registered Address over the Backbone using NA messages, asynchronously or as a response to a Neighbor Solicitation messages.
- o Deliver packets arriving from the LLN, using Neighbor Solicitation messages to look up the destination over the Backbone.
- o Forward or bridge packets between the LLN and the Backbone.
- o 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 connecting the Multi-link IPv6 subnet to the Internet.

The proxy-ND operation can co-exist with IPv6 ND over the Backbone.

The 6BBR may co-exist with a proprietary snooping or a traditional bridging functionality in an Access Point, in order to support legacy nodes that do not support this specification. In the case, the co-existing function may turn multicasts into a series of unicast to the legacy nodes.

The registration to a proxy service uses an NS/NA(EARO) exchange. 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 use the Extended Address Registration Option (EARO) defined in [RFC8505] as follows:

- o The EARO 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 with a 6LBR, if one is present, and the 6BBR. Address duplication is detected using the ROVR field. Conflicting registrations to different 6BBRs for the same Registered Address are resolved using the TID field.
- o 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. [RFC8505] requires that the registration NS(EARO) contains an SLLAO. This specification details the use of those messages over the backbone.

Note: [RFC6775] requires that the registration NS(EARO) contains an SLLAO and [RFC4862] 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 adds the capability to insert IPv6 ND options in the EDAR and EDAC messages. In particular, 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].

3.2. Access Link

Figure 2 illustrates a flow where 6LN forms an IPv6 Address and registers it to a 6BBR acting as a 6LR [RFC8505]. The 6BBRs applies ODAD (see Section 3.6) to the registered address to enable connectivity while the message flow is still in progress. In that 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.




Figure 2: Initial Registration Flow to a 6BBR acting as Routing Proxy

3.3. Route-Over Mesh

Figure 3 illustrates IPv6 signaling that enables a 6LN to form a Global or a Unique-Local Address and register it to the 6LBR that serves its LLN using [RFC8505]. The 6LBR (acting as Registering Node) proxies the registration to the 6BBR, using [RFC8505] to register the addresses the 6LN (Registered Node) on its behalf to the 6BBR, and obtain proxy-ND services from the 6BBR.




Figure 3: 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 RPL [RFC6550] and collocating the RPL root with a 6LBR that serves the LLN, and is either collocated with or connected to the 6BBR over an IPv6 Link.

3.4. The Binding Table

Addresses in a 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]. A 6BBR maintains a state for its active registrations 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:

- o De-registrations (newer TID, same ROVR, null Lifetime) are accepted with a status of 4 ("Removed"); the entry is deleted;
- o 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.
- o Identical registrations (same TID, same ROVR) from a 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 MUST be eventually produced, carrying the result of the DAD process;
- o Older registrations (older TID, same ROVR) from the same Registering Node are discarded;
- o 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;
- o Any registration for the same address but with a different ROVR is rejected with a status of 1 (Duplicate).

3.5. Primary and Secondary 6BBRs

A same address may be successfully registered to more than one 6BBR, in which case the Registering Node uses the same EARO in all the parallel registrations. To allow for this, ND(DAD) and NA messages with an EARO that indicate an identical Binding in another 6BBR (same Registered address, same TID, same ROVR) as silently ignored.

A 6BBR MAY optionally be 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.

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 2, 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 MUST be 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 (O) flag [RFC4861] that is not set. The router can then determine the freshest EARO in case of a conflicting NA(EARO) messages, using the method described in section 5.2.1 of [RFC8505]. If the NA(EARO) is the freshest answer, the default router creates a Binding with the SLLAO of the 6BBR (in Routing Proxy mode) or that of the Registering Node (in Bridging Proxy mode) so that traffic from/to the Registered Address can flow immediately.

4. MultiLink Subnet Considerations

The Backbone and the federated LLN Links are considered as different links in the MultiLink 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).

Nodes located inside the subnet do not perform the IPv6 Path MTU Discovery [RFC8201]. For that reason, the MTU must have a same value on the Backbone and all attached LLNs. To achieve this, 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 MultiLink 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 for duplication or 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.

This specification enables an address to be registered to more than one 6BBR. It results that a 6LBR MUST be capable to maintain a state for each of the 6BBR having registered with a same TID and same ROVR.

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 EDAC message SHOULD carry the SLLAO used in NS messages by the 6BBR for that Binding, and the EDAR message SHOULD carry the 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, respectively.

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.

An NA message generated in response to an NS(DAD) MUST have the Override flag set and a status of 1 (Duplicate) or 3 (Moved) in the EARO. An NA message generated in response to an NS(Lookup) or an NS(NUD) MUST NOT have the Override flag set.

This specification enables proxy operation for the IPv6 ND resolution of LLN devices and a prefix that is used across a MultiLink 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 Neighbor Unreachability Detection (NUD) 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, the MLSN prefix 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 a new 6BBR, the previous 6BBR SHOULD unicast a gratuitous NA with the Override flag set 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 with the Override flag; this NA will possibly hit all the nodes on the Backbone, whether or not they maintain an NCE for the Registered Address.

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 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 MultiLink 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 directly to the NS(Lookup) message or reinject the message on the backbone, either as a Layer-2 unicast or a multicast.

If the Registering Node owns the Registered Address, then its mobility does not impact existing NCEs over the Backbone. Otherwise, when the 6LN selects another Registering Node, the new Registering Node SHOULD send a multicast NA with the Override flag set to fix the existing NCEs across the Backbone. This method can fail if the multicast message is not received; one or more correspondent nodes on the Backbone might maintain a stale NCE, and packets to the Registered Address may be lost. When this condition happens, it is eventually be discovered and resolved using Neighbor Unreachability Detection (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 on the registering node in the LLN interface from which the registration was received.

The 6LR operation is modified as follows:

- o EDAR and EDAC messages SHOULD carry a SLLAO and a TLLAO, respectively.
- o A Bridging Proxy MAY register Link Local addresses to the 6BBR and proxy ND for those addresses over the backbone.
- o 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, and unless

administratively overridden, 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, 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 as it was. 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 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. It MAY preserve a temporary state in order to forward packets in flight. The state may be a NCE formed based on a received NA message, or a Binding in Stale state and pointing at the new 6BBR on the backbone.

The implementation should also use REDIRECT messages as specified in [RFC4861] to update the correspondents for the Registered Address, pointing the new 6BBR.

9.1. Operation 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:

- o 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 with a status of 1 (Duplicate) that indicates an existing registration owned by a different Registering Node. In that case, an NA MUST be sent back to the Registering Node with a status of 1 (Duplicate) in the EARO. 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]).

- o 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. This behavior might be overridden by policy, in particular if the registration is not trusted.
- o The Binding MUST be removed if an NA message is received over the Backbone for the Registered Address containing an EARO with a status of 3 (Moved), or an NS(DAD) with an EARO that indicates a fresher registration ([RFC8505]) for the same Registered Node (i.e. same ROVR). A status of 3 is returned in the NA(EARO) back to the Registering Node.
- o NS(DAD) and NA messages containing an EARO that indicates a registration for the same Registered Node that is not as fresh as this SHOULD be answered with an NA message containing an EARO with a status of 3 (Moved) in order to clean up the situation immediately.
- o Other NS(DAD) and NA messages from the Backbone are ignored.
- o 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 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 set over the backbone (see Section 7 and Section 8).

9.2. Operation on a Binding in Reachable State

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

An NS(DAD) with no EARO or with an EARO that indicates a duplicate 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:

- o 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). 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 small timer in order to allow for a parallel registration to arrive to this node, in which case the NA might be ignored.
- o 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.
- o 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).
- o Other NS(DAD) and NA messages from the Backbone are ignored.
- o 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.

9.3. Operation 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:

- o 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.
- o 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).
- o 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.
- o Other NS(DAD) and NA messages from the Backbone are ignored.

When the STALE_DURATION 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.

The Registering Node SHOULD 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 if the 6BBR cancels the NS(Lookup) over the LLN or to selected LLN nodes that are known to 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 [RFC7772] in order to limit the use of multicast RAs. It SHOULD also implement Simple Procedures for Detecting Network Attachment in IPv6 [RFC6059] (DNA procedures) to detect movements, and support Packet-Loss Resiliency for Router Solicitations [RFC7559] in order to improve reliability for the unicast RS messages.

11. Security Considerations

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

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.

[I-D.ietf-6lo-ap-nd] guarantees the ownership of a registered address based on a proof-of-ownership encoded in the ROVR field and protects against address theft and impersonation.

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, by default 24 hours. In LLNs where addresses are renewed rapidly, e.g. for privacy reasons, STALE_DURATION SHOULD be configured with a relatively long 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.

15. References

15.1. 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, <<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>>.
- [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>>.
- [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>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.
- [RFC8201] McCann, J., Deering, S., Mogul, J., and R. Hinden, Ed., "Path MTU Discovery for IP version 6", STD 87, RFC 8201, DOI 10.17487/RFC8201, July 2017, <<https://www.rfc-editor.org/info/rfc8201>>.
- [RFC8505] Thubert, P., Ed., Nordmark, E., Chakrabarti, S., and C. Perkins, "Registration Extensions for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Neighbor Discovery", RFC 8505, DOI 10.17487/RFC8505, November 2018, <<https://www.rfc-editor.org/info/rfc8505>>.

15.2. Informative References

- [I-D.bi-savi-wlan]
Bi, J., Wu, J., Wang, Y., and T. Lin, "A SAVI Solution for WLAN", draft-bi-savi-wlan-16 (work in progress), November 2018.
- [I-D.ietf-6lo-ap-nd]
Thubert, P., Sethi, M., Struik, R., and B. Sarikaya, "Address Protected Neighbor Discovery for Low-power and Lossy Networks", draft-ietf-6lo-ap-nd-09 (work in progress), December 2018.
- [I-D.ietf-6man-rs-refresh]
Nordmark, E., Yourtchenko, A., and S. Krishnan, "IPv6 Neighbor Discovery Optional RS/RA Refresh", draft-ietf-6man-rs-refresh-02 (work in progress), October 2016.
- [I-D.ietf-6tisch-architecture]
Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", draft-ietf-6tisch-architecture-19 (work in progress), December 2018.

[I-D.ietf-mboned-ieee802-mcast-problems]

Perkins, C., McBride, M., Stanley, D., Kumari, W., and J. Zuniga, "Multicast Considerations over IEEE 802 Wireless Media", draft-ietf-mboned-ieee802-mcast-problems-04 (work in progress), November 2018.

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

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

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

Thubert, P. and E. Levy-Abegnoli, "IPv6 Neighbor Discovery Unicast Lookup", draft-thubert-6lo-unicast-lookup-00 (work in progress), January 2019.

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

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

[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)".
- [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>>.
- [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>>.
- [RFC6606] Kim, E., Kaspar, D., Gomez, C., and C. Bormann, "Problem Statement and Requirements for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing", RFC 6606, DOI 10.17487/RFC6606, May 2012, <<https://www.rfc-editor.org/info/rfc6606>>.
- [RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", RFC 6830, DOI 10.17487/RFC6830, January 2013, <<https://www.rfc-editor.org/info/rfc6830>>.
- [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>>.
- [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>>.

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 trust 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, or FMIP, 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 MultiLink 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, meeting the requirements listed in Appendix B.3 of [RFC8505] "Requirements Related to Various Low-Power Link Types".

Each LLN in the subnet is attached at 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. In 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 IPv6 ND operation is minimized as the number of 6LNs grows in the LLN. 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 MultiLink 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
MOUGINS - Sophia Antipolis 06254
FRANCE

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

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

Email: charliep@computer.org

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

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

6Lo Working Group
Internet-Draft

Intended status: Standards Track
Expires: September 10, 2019

C. Gomez
S. Darroudi
Universitat Politecnica de Catalunya
T. Savolainen
DarkMatter
M. Spoerk
Graz University of Technology
March 9, 2019

IPv6 Mesh over BLUETOOTH(R) Low Energy using IPSP
draft-ietf-6lo-blemesh-05

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 mechanisms that are needed to enable IPv6 mesh over Bluetooth Low Energy links established by using the Bluetooth Internet Protocol Support Profile. This document does not specify the routing protocol to be used in an IPv6 mesh over Bluetooth LE links.

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 September 10, 2019.

Copyright Notice

Copyright (c) 2019 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 links	4
3.1. Protocol stack	4
3.2. Subnet model	5
3.3. Link model	6
3.3.1. Stateless address autoconfiguration	6
3.3.2. Neighbor Discovery	6
3.3.3. Header compression	7
3.3.4. Unicast and multicast mapping	8
4. IANA Considerations	9
5. Security Considerations	9
6. Contributors	9
7. Acknowledgements	9
8. Appendix	10
9. References	13
9.1. Normative References	13
9.2. Informative References	14
Authors' Addresses	14

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, the functionality described in RFC 7668 is not sufficient and would fail to enable an IPv6 mesh over Bluetooth LE

links. This document specifies mechanisms that are needed to enable IPv6 mesh over Bluetooth LE links. This document does not specify the routing protocol to be used in an IPv6 mesh over Bluetooth LE links.

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 (now deprecated) 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 and subsequent Bluetooth versions (e.g. Bluetooth 4.2 [BTCrev4.2] or subsequent), a device may implement both roles simultaneously.

This document assumes a mesh network composed of Bluetooth LE links, where link layer connections are established between neighboring IPv6-enabled devices (see Section 3.3.2, item 3.b)). 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 mesh over 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 links

3.1. Protocol stack

Figure 1 illustrates the protocol stack for IPv6 mesh over Bluetooth LE links. 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 over Bluetooth LE") is now adapted for IPv6 mesh over Bluetooth LE links, and b) the protocol stack for IPv6 mesh over Bluetooth LE links includes IPv6 routing functionality.




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

Bluetooth 4.2 defines a default MTU for Bluetooth LE of 251 bytes. Excluding the L2CAP header of 4 bytes, a protocol data unit (PDU) size of 247 bytes is available for the layer above L2CAP. (Note: earlier Bluetooth LE versions offered a maximum amount of 23 bytes for the layer atop L2CAP.) The L2CAP provides a fragmentation and reassembly solution for transmitting or receiving larger PDUs. At each link, the IPSP defines means for negotiating a link-layer connection that provides an MTU of 1280 octets or higher for the IPv6 layer [IPSP]. The link-layer MTU is negotiated separately for each direction. Implementations that require an equal link-layer MTU for the two directions SHALL use the smallest of the possibly different MTU values.

Note that this specification allows using different MTUs in different links. If an implementation requires use of the same MTU on every one of its links, and a new node with a smaller MTU is added to the network, a renegotiation of one or more links can occur. In the

worst case, the renegotiations could cascade network-wide. In that case, implementers need to assess the impact of such phenomenon.

Similarly to RFC 7668, fragmentation functionality from 6LoWPAN standards is not used for IPv6 mesh over Bluetooth LE links. Bluetooth LE's fragmentation support provided by L2CAP is used when necessary.

3.2. Subnet model

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

3.3. Link model

3.3.1. Stateless address autoconfiguration

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

Multihop DAD functionality as defined in section 8.2 of RFC 6775 and updated by RFC 8505, 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], subsequently updated by 'Registration Extensions for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Neighbor Discovery' [RFC8505], describes the neighbor discovery functionality adapted for use in several 6LoWPAN topologies, including the mesh topology. The route-over functionality of RFC 6775 and RFC 8505 MUST be supported.

The following aspects of the Neighbor Discovery optimizations for 6LoWPAN [RFC6775],[RFC8505] are applicable to Bluetooth LE 6LNs:

1. A Bluetooth LE host MUST register its non-link-local addresses with its routers by sending a Neighbor Solicitation (NS) message with the Extended Address Registration Option (EARO) and process the Neighbor Advertisement (NA) accordingly. The NS with the EARO option MUST be sent irrespective of the method used to generate the IID. The EARO option includes a Registration Ownership Verifier (ROVR) field [RFC8505]. In the case of Bluetooth LE, by default the ROVR field is filled with the 48-bit device address used by the Bluetooth LE node converted into 64-bit Modified EUI-64 format [RFC4291]. Optionally, a cryptographic ID (see [I-D.ietf-6lo-ap-nd]) MAY be placed in the ROVR field. If a cryptographic ID is used, address registration and multihop DAD formats and procedures defined in [I-D.ietf-6lo-ap-nd] MUST be used, unless an alternative mechanism offering equivalent protection is used. As per RFC 8505, a 6LN MUST NOT register its link-local address.

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 [RFC6775], and Section 5.6 of [RFC8505].

3. The router behavior for 6LRs and 6LBRs is described in Section 6 of RFC 6775, and updated by RFC 8505. However, as per this specification: a) Routers SHALL NOT use multicast NSs to discover other routers' link layer addresses. b) As per section 6.2 of RFC 6775, in a dynamic configuration scenario, a 6LR comes up as a non-router and waits to receive a Router Advertisement for configuring its own interface address first, before setting its interfaces to be advertising interfaces and turning into a router. In order to support such operation in an IPv6 mesh over Bluetooth LE links, a 6LR first uses the IPSP Node role only. Once the 6LR has established a connection with another node previously running as a router, and receives a Router Advertisement from that router, the 6LR configures its own interface address, it turns into a router, and it runs as an IPSP Router. A 6LBR uses the IPSP Router role since the 6LBR is initialized. See an example in the Appendix.

4. Border router behavior is described in Section 7 of RFC 6775, and updated by RFC 8505.

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). RFC 8505 updates those mechanisms and the related message formats. Implementations of this specification MUST support the features described in sections 8.1 and 8.2 of RFC 6775, as updated by RFC 8505, unless some alternative ("substitute") from some other specification is supported by the implementation.

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 an IPv6 mesh over 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).

A 6LN SHOULD register its non-link-local address with ARO in the next-hop router. Note that in some cases (e.g. very short-lived connections) it may not be worthwhile for a 6LN to send an NS with ARO for registering its address. 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 over Bluetooth LE links requires 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 links, 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.

The ROVR can be derived from the Bluetooth device address. However, such a ROVR can be spoofed, and therefore, any node connected to the subnet and aware of a registered-address-to-ROVR mapping could perform address theft and impersonation attacks. Use of Address Protected Neighbor Discovery [I-D.ietf-6lo-ap-nd] provides protection against such attacks.

6. Contributors

Carlo Alberto Boano (Graz University of Technology) contributed to the design and validation of this document.

7. 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.

6LN (not initialized) 6LN (not initialized) 6LN (not initialized)

Step 2

6LBR
(IPSP: Router)

6LR

(IPSP: Node)

6LR

(IPSP: Node)

6LN (not initialized) 6LN (not initialized) 6LN (not initialized)

Step 3

```

Bluetooth LE connection -->
                                6LBR
                                (IPSP: Router)
                               /      \
                              6LR      6LR
                              (IPSP: Node)  (IPSP: Node)

```

6LN (not initialized) 6LN (not initialized) 6LN (not initialized)

Step 4

```

6LBR
(IPSP: Router)
/          \

```

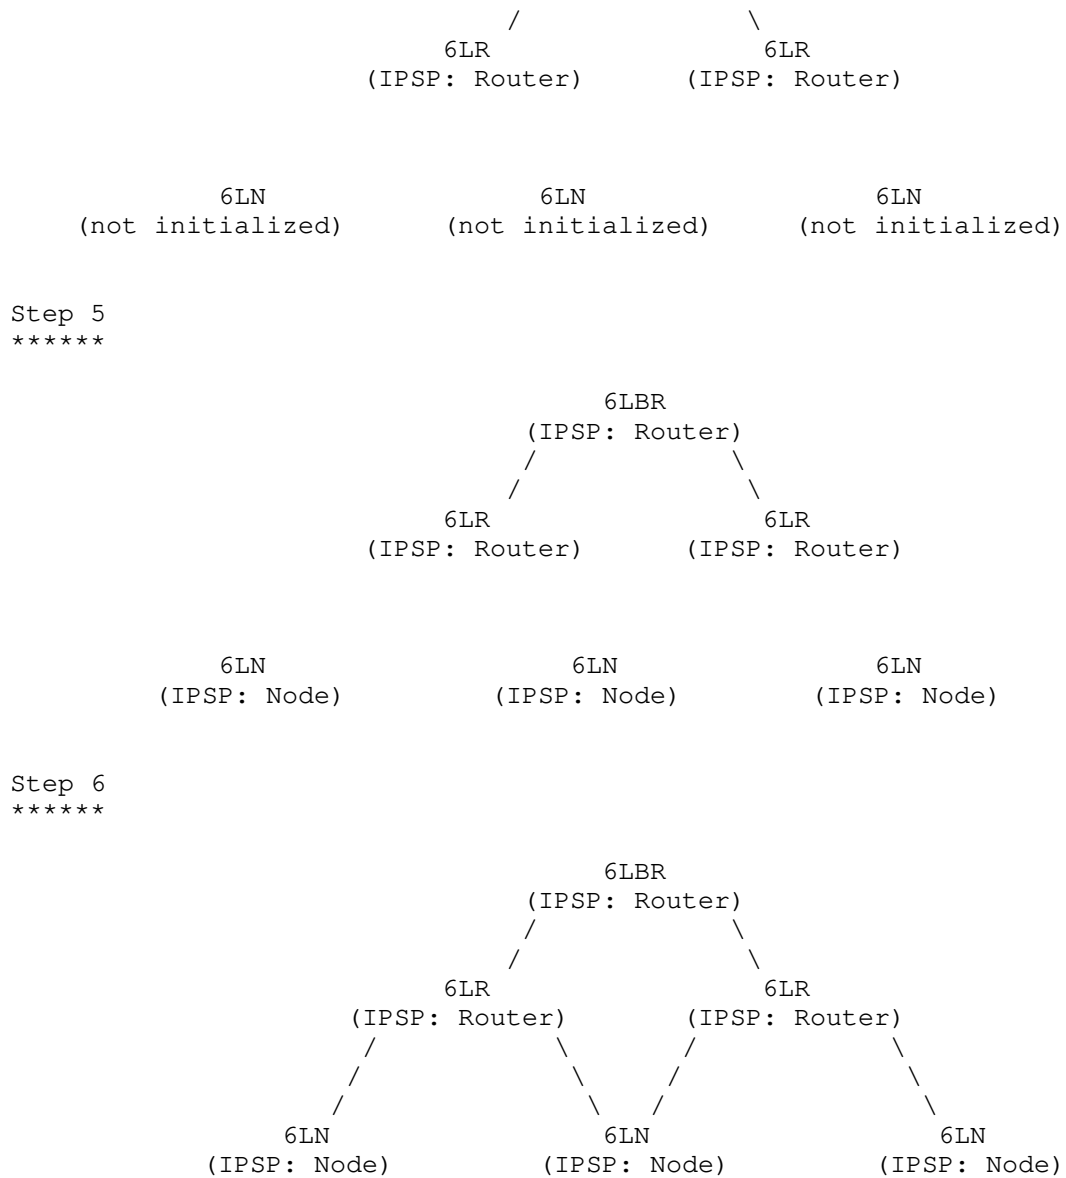


Figure 3: An example of connection establishment and use of IPSP roles in an IPv6 mesh over Bluetooth LE links.

9. References

9.1. Normative References

- [BTCorev4.2] Bluetooth Special Interest Group, "Bluetooth Core Specification Version 4.2", December 2014, <<https://www.bluetooth.com/specifications/archived-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>>.

- [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>>.

9.2. Informative References

- [BTCorev4.1] Bluetooth Special Interest Group, "Bluetooth Core Specification Version 4.1", December 2013, <<https://www.bluetooth.org/en-us/specification/adopted-specifications>>.
- [I-D.ietf-6lo-ap-nd] Thubert, P., Sarikaya, B., Sethi, M., and R. Struik, "Address Protected Neighbor Discovery for Low-power and Lossy Networks", draft-ietf-6lo-ap-nd-11 (work in progress), February 2019.
- [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
C/Esteve Terradas, 7
Castelldefels 08860
Spain

Email: carlesgo@entel.upc.edu

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

Email: sm.darroudi@entel.upc.edu

Teemu Savolainen
DarkMatter LLC

Email: teemu.savolainen@darkmatter.ae

Michael Spoerk
Graz University of Technology
Inffeldgasse 16/I
Graz 8010
Austria

Email: michael.spoerk@tugraz.at

6lo
Internet-Draft
Intended status: Standards Track
Expires: September 9, 2019

Lijo Thomas
C-DAC
S. Anamalamudi
SRM University-AP
S.V.R.Anand
Malati Hegde
Indian Institute of Science
C. Perkins
Futurewei
March 8, 2019

Packet Delivery Deadline time in 6LoWPAN Routing Header
draft-ietf-6lo-deadline-time-04

Abstract

This document specifies a new type for the 6LoWPAN routing header containing the deadline time for data packets, designed for use over constrained networks. The deadline time enables forwarding and scheduling decisions for time critical IoT M2M applications that operate within time-synchronized networks that agree on the meaning of the time representations used for the deadline time values.

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 September 9, 2019.

Copyright Notice

Copyright (c) 2019 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	3
3. 6LoRHE Generic Format	3
4. Deadline-6LoRHE	4
5. Deadline-6LoRHE Format	6
6. Deadline-6LoRHE in Three Network Scenarios	7
6.1. Scenario 1: Endpoints in the same DODAG (N1)	8
6.2. Scenario 2: Endpoints in Networks with Dissimilar L2 Technologies.	9
6.3. Scenario 3: Packet transmission across different DODAGs (N1 to N2).	10
7. IANA Considerations	11
8. Synchronization Aspects	12
9. Security Considerations	13
10. Acknowledgements	13
11. References	13
11.1. Normative References	13
11.2. Informative References	15
Appendix A. Changes from revision 03 to revision 04	16
Appendix B. Changes from revision 03 to revision 04	16
Appendix C. Changes from revision 01 to revision 02	17
Appendix D. Changes between earlier versions	17
Authors' Addresses	18

1. Introduction

Low Power and Lossy Networks (LLNs) are likely to be deployed for real time industrial applications requiring end-to-end delay guarantees [I-D.ietf-detnet-use-cases]. A Deterministic Network ("detnet") typically requires some data packets to reach their receivers within strict time bounds. Intermediate nodes use the deadline information to make appropriate packet forwarding and scheduling decisions to meet the time bounds.

This document specifies a new type for the Elective 6LoWPAN Routing Header (6LoRHE), so that the deadline time (i.e., the time of latest acceptable delivery) of data packets can be included within the 6LoWPAN routing header. [RFC8138] specifies the 6LoWPAN Routing

Header (6LoRH), compression schemes for RPL routing (source routing) operation [RFC6554], header compression of RPL Packet Information [RFC6553], and IP-in-IP encapsulation. This document also specifies handling of the deadline time when packets traverse between time-synchronized networks operating in different timezones or distinct reference clocks. Time synchronization techniques are outside the scope of this document. There are a number of standards available for this purpose, including IEEE 1588 [ieee-1588], IEEE 802.1AS [dot1AS-2011], IEEE 802.15.4-2015 TSCH [dot15-tsch], and more.

The Deadline-6LoRHE can be used in any time synchronized 6Lo network. A 6TiSCH network is used to describe the implementation of the Deadline-6LoRHE, but this does not preclude its use in scenarios other than 6TiSCH. For instance, there is a growing interest in using 6Lo over a BLE mesh network [I-D.ietf-6lo-blemesh] in industrial IoT [dotBLEMesh]. BLE mesh time synchronization is being explored by the Bluetooth community. There are also cases under consideration in Wi-SUN [Wi-SUN_PHY], [dotWi-SUN].

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119] [RFC8174].

This document uses the terminology defined in [RFC6550] and [I-D.ietf-6tisch-terminology].

3. 6LoRHE Generic Format

Note: this section is not normative and is included for convenience. The generic header format of the 6LoRHE is specified in [I-D.ietf-roll-routing-dispatch]. Figure 1 illustrates the 6LoRHE generic format.




Figure 1: 6LoRHE format

- o Length: Length of the 6LoRHE expressed in bytes, excluding the first 2 bytes. This enables a node to skip a 6LoRHE if the Type is not recognized/supported.

- o Type (variable length): Type of the 6LoRHE (see Section 7)

4. Deadline-6LoRHE

The Deadline-6LoRHE (see Figure 3) is an elective 6LoRH (i.e., a 6LoRHE [RFC8138]) that provides the Deadline Time (DT) for an IPv6 datagram in a compressed form. Along with the deadline, the header can include the packet Origination Time Delta (OTD), the time at which the packet is enqueued for transmission (expressed as a value to be subtracted from DT); this enables a close estimate of the total delay incurred by a packet. The OTD field is initialized by the sender based on the current time at the outgoing network interface through which the packet is forwarded. Since the OTD is a delta the length of the OTD field (i.e., OTL) will require fewer bits than the length of the DT field (i.e., DTL).

The deadline field contains the value of the deadline time for the packet. The packet SHOULD be delivered to the Receiver before this time.

$$\text{packet_deadline_time} = \text{packet_origination_time} + \text{max_delay}$$

All nodes within the network SHOULD process the Deadline-6LoRHE in order to support delay-sensitive deterministic applications. The packet deadline time (DT) and origination time (OTD) are represented in time units determined by a scaling parameter in the routing header. One of the time units is the Network ASN (Absolute Slot Number) which can be used in case of a time slotted synchronized network (for instance a 6TiSCH network, where global time is maintained in the units of slot lengths of a certain resolution).

The delay experienced by packets in the network is a useful metric for network diagnostics and performance monitoring. Whenever a packet crosses into a network using a different reference clock, the Destination Time field is updated to represent the same Destination Time, but expressed using the reference clock of the interface into the new network. Then the origination time is the same as the current time when the packet is transmitted into the new network, minus the delay already experienced by the packet, say 'dly'. In this way, within the newly entered network, the packet will appear to have originated 'dly' time units earlier with respect to the reference clock of the new network.

$$\text{origination time in new network} = \text{current_time_in_new_network} - \text{delay_already_experienced_in_previous_network(s)}$$

The following example illustrates these calculations when a packet travels between three networks, each in a different time zone. 'x'

can be 1, 2 or 3. Suppose that the deadline time as measured in timezone 1 is 1050 and the origination time is 50. Suppose that the difference between TZ2 and TZ1 is 900, and the the difference between TZ3 and TZ3 is 3600. In the figure, OT is the origination time as measured in the current timezone, and is equal to DT - OTD, that is, DT - 1000. Figure 2 uses the following abbreviations:

TxA : Time of arrival of packet in the network 'x'

TxD : Departure time of packet from the network 'x'

dlyx : Delay experienced by the packet in the previous network(s)

TZx : The time zone of network 'x'

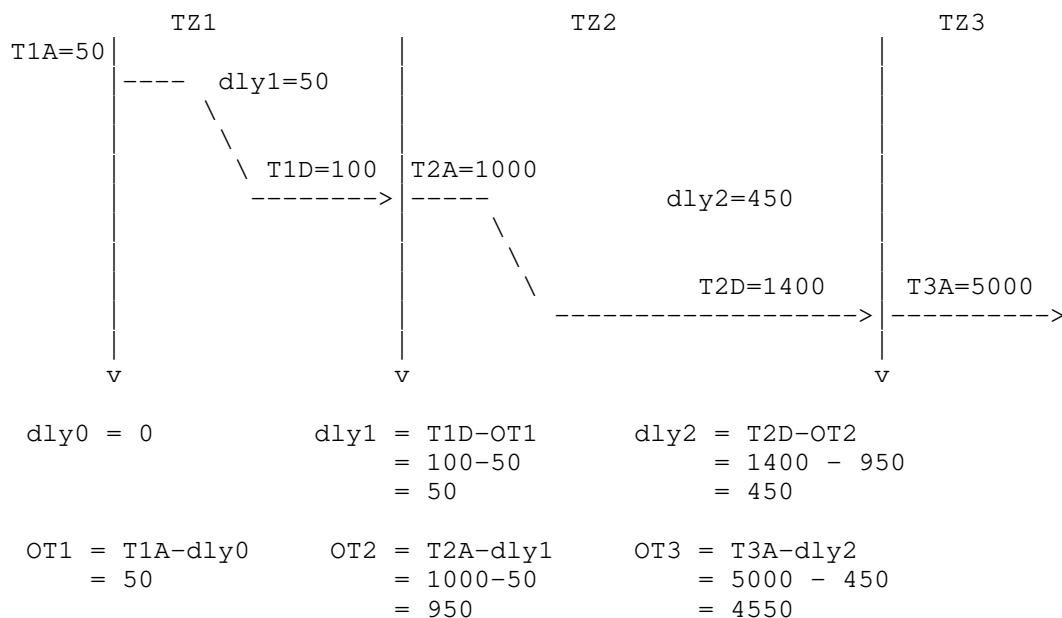


Figure 2: Destination Time Update example

There are multiple ways that a packet can be delayed, including queuing delay, MAC layer contention delay, serialization delay, and propagation delays. Sometimes there are processing delays as well. For the purpose of determining whether or not the deadline has already passed, these various delays are not distinguished.

5. Deadline-6LoRHE Format




Figure 3: Deadline-6LoRHE format

- o Length (5 bits): Length represents the total length of the Deadline-6LoRHE type measured in octets.
- o 6LoRH Type: TBD (see Section 7)
- o D flag (1 bit): The 'D' flag, set by the Sender, qualifies the action to be taken when a 6LR detects that the deadline time has elapsed. If 'D' bit is 1, then the 6LR MUST drop the packet if the deadline time is elapsed. If 'D' bit is 0, the packet MAY be forwarded on an exception basis, if the forwarding node is NOT in a situation of constrained resource, and if there are reasons to suspect that downstream nodes might find it useful (delay measurements, interpolations, etc.).
- o DTL (4 bits): Length of DT field as an unsigned 4-bit integer, encoding the length of the field in hex digits, minus one.
- o OTL (3 bits) : Length of OTD field as an unsigned 3-bit integer, encoding the length of the field in hex digits. If OTL == 0, the OTD field is not present. The value of OTL MUST NOT exceed the value of DTL plus one.
 - * For example, DTL = 0b0000 means the deadline time in the 6LoRHE is 1 hex digit (4 bits) long. OTL = 0b111 means the origination time is 7 hex digits (28 bits) long.
- o TU (2 bits) : Indicates the time units for DT and OTD fields. The encoding for the DT and OTD fields MUST always use the same time units and precision.
 - * 00 : Time represented in seconds and fractional seconds
 - * 01 : Reserved
 - * 10 : Network ASN
 - * 11 : Reserved
- o Binary Pt (6 bits) : If zero, the number of bits of the integer part the DT is equal to the number of bits of the fractional part of the DT. if nonzero, the Binary Pt is a signed integer determining the position of the binary point within the value for the DT.

- * If BinaryPt value is positive, then the number of bits for the integer part of the DT is increased by the value of BinaryPt, and the number of bits for the fractional part of the DT is correspondingly reduced. This increases the range of DT.
- * If BinaryPt value is negative, then the number of bits for the integer part of the DT is decreased by the value of BinaryPt, and the number of bits for the fractional part of the DT is correspondingly increased. This increases the precision of the fractional seconds part of DT.
- o DT Value (8..64-bit) : An unsigned integer of DTL+1 hex digits giving the Deadline Time value
- o OTD Value (8..64-bit) : An unsigned integer of OTL hex digits giving the Origination Time as a negative offset from the DT value

Whenever a sender initiates the IP datagram, it includes the Deadline-6LoRHE along with other 6LoRH information. For information about the time synchronization requirements between sender and receiver see Section 8.

Example: Consider a 6TiSCH network with time-slot length of 10ms. Let the time units be ASNs (TU == (binary)0b10). Let the current ASN when the packet is originated be 54400, and the maximum allowable delay (max_delay) for the packet delivery be 1 second from the packet origination, then:

```
deadline_time = packet_origination_time + max_delay
               = 0xD480 + 0x64 (Network ASNs)
               = 0xD4E4 (Network ASNs)
```

Then, the Deadline-6LoRHE encoding with nonzero OTL is:

```
DTL = 3, OTL = 2, TU = 0b10, BinaryPt = 8, DT = 0xD4E4, OTD
= 0x64
```

6. Deadline-6LoRHE in Three Network Scenarios

In this section, Deadline-6LoRHE operation is described for 3 network scenarios. Figure 4 depicts a constrained time-synchronized LLN that has two subnets N1 and N2, connected through LBRs [I-D.ietf-6lo-backbone-router] with different reference clock times T1 and T2.

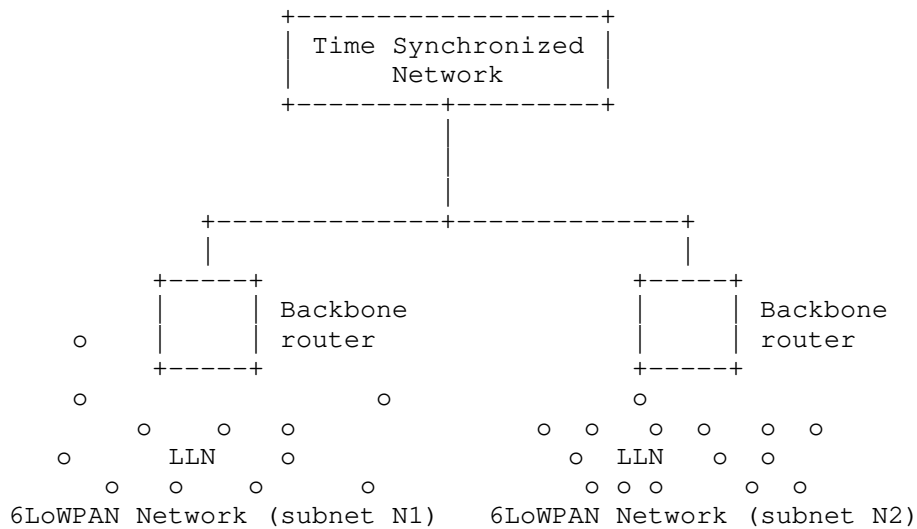


Figure 4: Intra-network Timezone Scenario

6.1. Scenario 1: Endpoints in the same DODAG (N1)

In scenario 1, shown in Figure 5, the Sender 'S' has an IP datagram to be routed to a Receiver 'R' within the same DODAG. For the route segment from Sender to 6LBR, the Sender includes a Deadline-6LoRHE by encoding the deadline time contained in the packet. Subsequently, each 6LR will perform hop-by-hop routing to forward the packet towards the 6LBR. Once 6LBR receives the IP datagram, it sends the packet downstream towards 'R'.

In case of a network running RPL non-storing mode, the 6LBR generates a IPv6-in-IPv6 encapsulated packet when sending the packet downwards to the Receiver [I-D.ietf-roll-useofrplinfo]. The 6LBR copies the Deadline-6LoRHE from the Sender originated IP header to the outer IP header. The Deadline-6LoRHE contained in the inner IP header is removed.




Figure 5: End points within same DODAG (subnet N1)

At the tunnel endpoint of the encapsulation, the Deadline-6LoRHE is copied back from the outer header to inner header, and the inner IP packet is delivered to 'R'.

6.2. Scenario 2: Endpoints in Networks with Dissimilar L2 Technologies.

In scenario 2, shown in Figure 6, the Sender 'S' (belonging to DODAG 1) has IP datagram to be routed to a Receiver 'R' over a time-synchronized IPv6 network. For the route segment from 'S' to 6LBR, 'S' includes a Deadline-6LoRHE. Subsequently, each 6LR will perform hop-by-hop routing to forward the packet towards the 6LBR. Once the Deadline Time information reaches the border router, the packet will be encoded according to the mechanism prescribed in the other time-synchronized network depicted as "Time Synchronized Network" in the figure 6. The specific data encapsulation mechanisms followed in the new network are beyond the scope of this document.

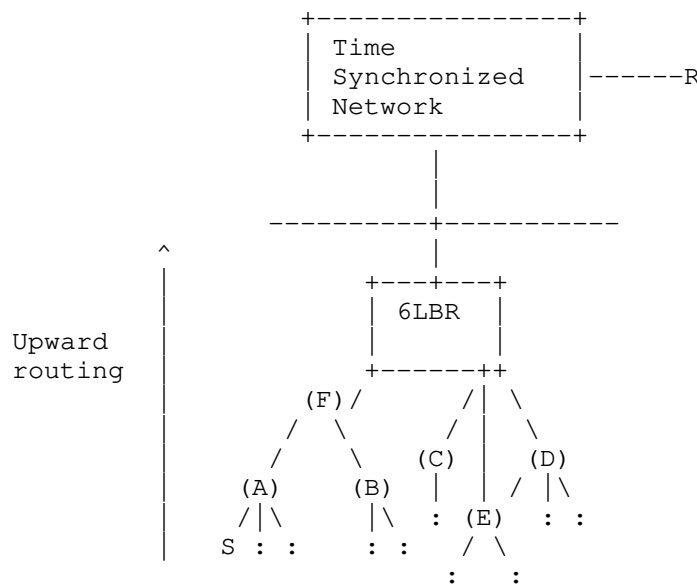


Figure 6: Packet transmission in Dissimilar L2 Technologies or Internet

For instance, the IP datagram could be routed to another time synchronized deterministic network using the mechanism specified in the In-band OAM [I-D.ietf-ippm-ioam-data], and then the deadline time would be updated according to the measurement of the current time in the new network.

6.3. Scenario 3: Packet transmission across different DODAGs (N1 to N2).

Consider the scenario depicted in Figure 7, in which the Sender 'S' (belonging to DODAG 1) has an IP datagram to be sent to Receiver 'R' belonging to another DODAG (DODAG 2). The operation of this scenario can be decomposed into combination of case 1 and case 2 scenarios. For the route segment from 'S' to 6LBR1, 'S' includes the Deadline-6LoRHE. Subsequently, each 6LR will perform hop-by-hop operation to forward the packet towards the 6LBR1. Once the IP datagram reaches 6LBR1 of DODAG1, it applies the same rule as described in Case 2 while routing the packet to 6LBR2 over a (likely) time synchronized wired backhaul. The wired side of 6LBR2 can be mapped to receiver of Case 2. Once the packet reaches 6LBR2, it updates the Deadline-6LoRHE by adding or subtracting the difference of time of DODAG2 and sends the packet downstream towards 'R'.




Figure 7: Packet transmission in different DODAGs (N1 to N2)

Consider an example of a 6TiSCH network in which S in DODAG1 generates the packet at ASN 20000 to R in DODAG2. Let the maximum allowable delay be 1 second. The time-slot length in DODAG1 and DODAG2 is assumed to be 10ms. Once the deadline time is encoded in Deadline-6LoRHE, the packet is forwarded to 6LBR of DODAG1. Suppose the packet reaches 6LBR of DODAG1 at ASN 20030.

```
current_time = ASN at LBR * slot_length_value

remaining_time = deadline_time - current_time
= ((packet_origination_time + max_delay) - current time)
= (20000 + 100) - 20030
= 30 (in Network ASNs)
= 30 * 10^3 milliseconds.
```

Once the Deadline Time information reaches the border router, the packet will be encoded according to the mechanism prescribed in the other time-synchronized network.

7. IANA Considerations

This document defines a new Elective 6LoWPAN Routing Header Type, and IANA is requested to assign a value (TBD) from the 6LoWPAN Dispatch Page1 number space for this purpose.

Elective 6LoRH Type	Value
Deadline-6LoRHE	TBD

Figure 8: Deadline-6LoRHE type

8. Synchronization Aspects

The document supports time representation of the deadline and origination times carried in the packets traversing through networks of different time zones having different time synchronization mechanisms. For instance, in a 6TiSCH network where the time is maintained as ASN time slots, the time synchronization is achieved through beaconing among the nodes as described in [RFC7554]. There could be 6lo networks that employ NTP where the nodes are synchronized with an external reference clock from an NTP server. The specification of the time synchronization method that need to be followed by a network is beyond the scope of the document.

The number of hex digits chosen to represent DT, and the portion of that field allocated to represent integer number of seconds, determines the meaning of t_0 , i.e., the meaning of $DT == 0$ in the chosen representation. If $DTL == 0$, then there are only 4 bits that can be used to count the time units, so that $DT == 0$ can never be more than 16 time units in the past. This then requires that the time synchronization between sender and receiver has to be tighter than 16 time units. If the binary point were moved so that all the bits were used for fractional time units (e.g., fractional seconds or fractional ASNs), the time synchronization requirement would be correspondingly tighter.

A 4-bit field for DT allows up to 16 hex digits, which is 64 bits. That is enough to represent the NTP [RFC5905] 64-bit timestamp format, which is more than enough for the purposes of establishing deadline times. Unless the binary point is moved, this is enough to represent time since year 1900.

For example, suppose that $DTL = 0b0000$ and the DT bits are split evenly; then we can count up to 3 integer seconds. In that case t_0 would be the most recent second of the current minute that has $t \bmod 4 == 0$. In other words, t_0 could be 0, 4, 8, 12, 16, ..., 52, or 56 seconds since the start of the most recent minute. The networks have to be synchronized well enough to ensure detection of overrun, and therefore to know which of those values is the correct value for t_0 . This is the hardest case.

If $DT = 3$ and the DT bits are again split evenly, then we can count up to 4,096 seconds. t_0 would be the start of the most recent hour.

For $TU = 0b00$, the time units are seconds. With $DTL == 15$, and $Binary\ Pt == 0$, the epoch is (by default) January 1, 1900 at 00:00 UTC. The resolution is then $(2^{(-32)})$ seconds, which is the maximum possible. This time format wraps around every 2^{32} seconds, which is roughly 136 years. For other choices of DTL and the $Binary\ Pt$, the value of t_0 (i.e., the meaning of $DT == 0$) needs to be established by means out of scope of this document.

For $TU = 0b10$, the time units are ASNs. The start time is relative, and updated by a mechanism out of scope for this document. With 10 ms slots, $DTL = 15$, and $Binary\ Pt == 0$, it would take over a year for the ASN to wrap around. Typically, the number of hex digits allocated for $TU = 0b10$ would be less than 15.

9. Security Considerations

The security considerations of [RFC4944], [RFC6282] and [RFC6553] apply. Using a compressed format as opposed to the full in-line format is logically equivalent and does not create an opening for a new threat when compared to [RFC6550], [RFC6553] and [RFC6554].

10. Acknowledgements

The authors thank Pascal Thubert for suggesting the idea and encouraging the work. Thanks to Shwetha Bhandari's suggestions which were instrumental in extending the timing information to heterogeneous networks. The authors acknowledge the 6TiSCH WG members for their inputs on the mailing list. Special thanks to Jerry Daniel, Seema Kumar, Avinash Mohan, Shalu Rajendran and Anita Varghese for their support and valuable feedback.

11. References

11.1. Normative References

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

Palattella, M., Thubert, P., Watteyne, T., and Q. Wang,
"Terms Used in IPv6 over the TSCH mode of IEEE 802.15.4e",
draft-ietf-6tisch-terminology-10 (work in progress), March
2018.

[I-D.ietf-roll-routing-dispatch]

Thubert, P., Bormann, C., Toutain, L., and R. Cragie,
"6LoWPAN Routing Header", draft-ietf-roll-routing-
dispatch-05 (work in progress), October 2016.

- [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>>.
- [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>>.
- [RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <<https://www.rfc-editor.org/info/rfc5905>>.
- [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>>.
- [RFC6553] Hui, J. and JP. Vasseur, "The Routing Protocol for Low-Power and Lossy Networks (RPL) Option for Carrying RPL Information in Data-Plane Datagrams", RFC 6553, DOI 10.17487/RFC6553, March 2012, <<https://www.rfc-editor.org/info/rfc6553>>.
- [RFC6554] Hui, J., Vasseur, JP., Culler, D., and V. Manral, "An IPv6 Routing Header for Source Routes with the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6554, DOI 10.17487/RFC6554, March 2012, <<https://www.rfc-editor.org/info/rfc6554>>.
- [RFC7554] Watteyne, T., Ed., Palattella, M., and L. Grieco, "Using IEEE 802.15.4e Time-Slotted Channel Hopping (TSCH) in the Internet of Things (IoT): Problem Statement", RFC 7554, DOI 10.17487/RFC7554, May 2015, <<https://www.rfc-editor.org/info/rfc7554>>.

- [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>>.
- [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>>.

11.2. Informative References

- [dot15-tsch]
"IEEE 802 Wireless", "IEEE Standard for Low-Rate Wireless Networks, Part 15.4, IEEE Std 802.15.4-2015", April 2016.
- [dot1AS-2011]
"IEEE Standards", "IEEE Standard for Local and Metropolitan Area Networks - Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks", March 2011.
- [dotBLEMesh]
Leonardi, L., Pattim, G., and L. Lo Bello, "Multi-Hop Real-Time Communications Over Bluetooth Low Energy Industrial Wireless Mesh Networks", IEEE Access Vol 6, 26505-26519, May 2018.
- [dotWi-SUN]
Harada, H., Mizutani, K., Fujiwara, J., Mochizuki, K., Obata, K., and R. Okumura, "IEEE 802.15.4g Based Wi-SUN Communication Systems", IEICE Transactions on Communications volume E100.B, Jan 2017.
- [I-D.ietf-6lo-backbone-router]
Thubert, P., Perkins, C., and E. Levy-Abegnoli, "IPv6 Backbone Router", draft-ietf-6lo-backbone-router-11 (work in progress), February 2019.
- [I-D.ietf-6lo-blemesh]
Gomez, C., Darroudi, S., Savolainen, T., and M. Spoerk, "IPv6 Mesh over BLUETOOTH(R) Low Energy using IPSP", draft-ietf-6lo-blemesh-04 (work in progress), January 2019.
- [I-D.ietf-detnet-use-cases]
Grossman, E., "Deterministic Networking Use Cases", draft-ietf-detnet-use-cases-20 (work in progress), December 2018.

[I-D.ietf-ippm-ioam-data]

Brockners, F., Bhandari, S., Pignataro, C., Gredler, H., Leddy, J., Youell, S., Mizrahi, T., Mozes, D., Lapukhov, P., Chang, R., daniel.bernier@bell.ca, d., and J. Lemon, "Data Fields for In-situ OAM", draft-ietf-ippm-ioam-data-04 (work in progress), October 2018.

[I-D.ietf-roll-useofrplinfo]

Robles, I., Richardson, M., and P. Thubert, "Using RPL Option Type, Routing Header for Source Routes and IPv6-in-IPv6 encapsulation in the RPL Data Plane", draft-ietf-roll-useofrplinfo-24 (work in progress), January 2019.

[ieee-1588]

"IEEE Standards", "IEEE Std 1588-2008 Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", July 2008.

[Wi-SUN_PHY]

Wi-SUN Alliance, "Wi-SUN PHY Specification V1.0", March 2016.

Appendix A. Changes from revision 03 to revision 04

This section lists the changes between draft-ietf-6lo-deadline-time revisions ...-03.txt and ...-04.txt.

- o Replaced OT (Origination Time) field by OTD (Origination Time Delta), allowing a more compressed representation that needs less processing during transitions between networks.
- o Changed representation for DTL, OTL, DT, OTD. Eliminated EXP in favor of BinaryPt.
- o Revised the figures and examples to use new parameters
- o Added new section on Synchronization Aspects to supply pertinent information about how nodes agree on the meaning of t=0.
- o Responded to numerous reviewer comments to improve editorial consistency and improve terminology.

Appendix B. Changes from revision 03 to revision 04

This section lists the changes between draft-ietf-6lo-deadline-time revisions ...-02.txt and ...-03.txt.

- o Added non-normative 6LoRHE description, citing RFC 8138.

- o Specified that the Origination Time (OT) is the time that packet is enqueued for transmission.
- o Mentioned more sources of packet delay.
- o Clarified reasons that packet MAY be forwarded if 'D' bit is 0.
- o Clarified that DT, OT, DTL and OTL are unsigned integers.
- o Updated bibliographic citations, including BLEmesh and Wi-SUN.

Appendix C. Changes from revision 01 to revision 02

This section lists the changes between draft-ietf-6lo-deadline-time revisions ...-01.txt and ...-02.txt.

- o Replaced 6LoRHE description by reference to RFC 8138.
- o Added figure to illustrate change to Origination Time when a packet crosses timezone boundaries.
- o Clarified that use of 6tisch networks is descriptive, not normative.
- o Clarified that In-Band OAM is used as an example and is not normative.
- o Updated bibliographic citations.
- o Alphabetized contributor names.

Appendix D. Changes between earlier versions

This section lists the changes between draft-ietf-6lo-deadline-time revisions ...-00.txt and ...-01.txt.

- o Changed "SHOULD drop" to "MUST drop" a packet if the deadline is passed (see Section 5).
- o Added explanatory text about how packet delays might arise. (see Section 4).
- o Mentioned availability of time-synchronization protocols (see Section 1).
- o Updated bibliographic citations.
- o Alphabetized contributor names.

- o Added this section.

Authors' Addresses

Lijo Thomas
C-DAC
Centre for Development of Advanced Computing (C-DAC), Vellayambalam
Trivandrum 695033
India

Email: lijo@cdac.in

Satish Anamalamudi
SRM University-AP
Amaravati Campus
Amaravati, Andhra Pradesh 522 502
India

Email: satishnaidu80@gmail.com

S.V.R Anand
Indian Institute of Science
Bangalore 560012
India

Email: anand@ece.iisc.ernet.in

Malati Hegde
Indian Institute of Science
Bangalore 560012
India

Email: malati@ece.iisc.ernet.in

Charles E. Perkins
Futurewei
2330 Central Expressway
Santa Clara 95050
Unites States

Email: charliep@computer.org

6lo
Internet-Draft
Updates: 4944 (if approved)
Intended status: Standards Track
Expires: July 27, 2019

P. Thubert, Ed.
Cisco Systems
January 23, 2019

6LoWPAN Selective Fragment Recovery
draft-ietf-6lo-fragment-recovery-02

Abstract

This draft updates RFC 4944 with a simple protocol to recover individual fragments across a route-over mesh network, with a minimal flow control to protect the network against bloat.

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 July 27, 2019.

Copyright Notice

Copyright (c) 2019 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	3
2.1. BCP 14	3
2.2. References	4
2.3. 6LoWPAN Acronyms	4
2.4. Referenced Work	4
2.5. New Terms	5
3. Updating RFC 4944	5
4. Updating draft-wattheyne-6lo-minimal-fragment	6
4.1. Slack in the First Fragment	6
4.2. Modifying the First Fragment	6
5. New Dispatch types and headers	7
5.1. Recoverable Fragment Dispatch type and Header	8
5.2. RFRAG Acknowledgment Dispatch type and Header	9
6. Fragments Recovery	11
7. Forwarding Fragments	13
7.1. Upon the first fragment	13
7.2. Upon the next fragments	13
7.3. Upon the RFRAG Acknowledgments	14
8. Security Considerations	14
9. IANA Considerations	15
10. Acknowledgments	15
11. References	15
11.1. Normative References	15
11.2. Informative References	16
Appendix A. Rationale	18
Appendix B. Requirements	19
Appendix C. Considerations On Flow Control	20
Author's Address	21

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 IEEE Std. 802.15.4 [IEEE.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 10Kbytes or more and an end-to-end reliable transport is required.

"Transmission of IPv6 Packets over IEEE 802.15.4 Networks" [RFC4944] defines the original 6LoWPAN datagram fragmentation mechanism for LLNs. One critical issue with this original design is that routing an IPv6 [RFC8200] packet across a route-over mesh requires to reassemble the full packet at each hop, which may cause latency along a path and an overall buffer bloat in the network. The "6TiSCH Architecture" [I-D.ietf-6tisch-architecture] recommends to use a hop-by-hop fragment forwarding technique to alleviate those undesirable effects. "LLN Minimal Fragment Forwarding" [I-D.wattheyne-6lo-minimal-fragment] proposes such a technique, in a fashion that is compatible with [RFC4944] without the need to define a new protocol. However, adding that capability alone to the local implementation of the original 6LoWPAN fragmentation would not address the bulk of the issues raised against it, and may create new issues like remnant state in the network.

Another issue against [RFC4944] is that it does not define a mechanism to first discover the loss of a fragment along a multi-hop path (e.g. having exhausted the link-layer retries at some hop on the way), and then to recover that loss. With RFC 4944, the forwarding of a whole datagram fails when one fragment is not delivered properly to the destination 6LoWPAN endpoint. End-to-end transport or application-level mechanisms may require a full retransmission of the datagram, wasting resources in an already constrained network.

In that situation, the source 6LoWPAN endpoint will not be aware that a loss occurred and will continue sending all fragments for a datagram that is already doomed. The original support is missing signaling to abort a multi-fragment transmission at any time and from either end, and, if the capability to forward fragments is implemented, clean up the related state in the network. It is also lacking flow control capabilities to avoid participating to a congestion that may in turn cause the loss of a fragment and trigger the retransmission of the full datagram.

This specification proposes a method to forward fragments across a multi-hop route-over mesh, and to recover individual fragments between LLN endpoints. The method is designed to limit congestion loss in the network and addresses the requirements that are detailed in Appendix B.

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. References

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

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

2.3. 6LoWPAN Acronyms

This document uses the following acronyms:

6BBR: 6LoWPAN Backbone Router

6LBR: 6LoWPAN Border Router

6LN: 6LoWPAN Node

6LR: 6LoWPAN Router

LLN: Low-Power and Lossy Network

2.4. Referenced Work

Past experience with fragmentation has shown that miss-associated or lost fragments can lead to poor network behavior and, occasionally, trouble at application layer. The reader is encouraged to read "IPv4 Reassembly Errors at High Data Rates" [RFC4963] and follow the references for more information.

That experience led to the definition of "Path MTU discovery" [RFC8201] (PMTUD) protocol that limits fragmentation over the Internet.

Specifically in the case of UDP, valuable additional information can be found in "UDP Usage Guidelines for Application Designers" [RFC8085].

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].

"The Benefits of Using Explicit Congestion Notification (ECN)" [RFC8087] provides useful information on the potential benefits and pitfalls of using ECN.

Quoting the "Multiprotocol Label Switching (MPLS) Architecture" [RFC3031]: with MPLS, "packets are "labeled" before they are forwarded. At subsequent hops, there is no further analysis of the packet's network layer header. Rather, the label is used as an index into a table which specifies the next hop, and a new label". The MPLS technique is leveraged in the present specification to forward fragments that actually do not have a network layer header, since the fragmentation occurs below IP.

"LLN Minimal Fragment Forwarding" [I-D.watteyne-6lo-minimal-fragment] introduces the concept of a Virtual Reassembly Buffer (VRB) and an associated technique to forward fragments as they come, using the datagram_tag as a label in a fashion similar to MLPS. This specification reuses that technique with slightly modified controls.

2.5. New Terms

This specification uses the following terms:

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. Updating RFC 4944

This specification updates the fragmentation mechanism that is specified in "Transmission of IPv6 Packets over IEEE 802.15.4 Networks" [RFC4944] for use in route-over LLNs by providing a model where fragments can be forwarded end-to-end across a 6LoWPAN LLN, and where fragments that are lost on the way can be recovered individually. A new format for fragment is introduces and new dispatch types are defined in Section 5.

[RFC8138] allows to modifies the size of a packet en-route by removing the consumed hops in a compressed Routing Header. It results that the fragment_offset and datagram_size cannot be signaled in the uncompressed form. This specification expresses those fields in the compressed form and allows to modify them en-route (see Section 4.2.

Note that consistantly with in Section 2 of [RFC6282] for the fragmentation mechanism described in Section 5.3 of [RFC4944], any

header that cannot fit within the first fragment MUST NOT be compressed when using the fragmentation mechanism described in this specification.

4. Updating draft-wattheyne-6lo-minimal-fragment

This specification updates the fragment forwarding mechanism specified in "LLN Minimal Fragment Forwarding" [I-D.wattheyne-6lo-minimal-fragment] by providing additional operations to improve the management of the Virtual Reassembly Buffer (VRB).

4.1. Slack in the First Fragment

At the time of this writing, [I-D.wattheyne-6lo-minimal-fragment] allows for refragmenting in intermediate nodes, meaning that some bytes from a given fragment may be left in the VRB to be added to the next fragment. The reason for this to happen would be the need for space in the outgoing fragment that was not needed in the incoming fragment, for instance because the 6LoWPAN Header Compression is not as efficient on the outgoing link, e.g., if the Interface ID (IID) of the source IPv6 address is elided by the originator on the first hop because it matches the source MAC address, but cannot be on the next hops because the source MAC address changes.

This specification cannot allow this operation since fragments are recovered end-to-end based on the fragment number. This means that the fragments that contain a 6LoWPAN-compressed header MUST have enough slack to enable a less efficient compression in the next hops that still fits in one MAC frame. For instance, if the IID of the source IPv6 address is elided by the originator, then it MUST compute the `fragment_size` as if the MTU was 8 bytes less. This way, the next hop can restore the source IID to the first fragment without impacting the second fragment.

4.2. Modifying the First Fragment

The compression of the Hop Limit, of the source and destination addresses, and of the Routing Header may change en route in a Route-Over mesh LLN. If the size of the first fragment is modified, then the intermediate node MUST adapt the `datagram_size` to reflect that difference.

The intermediate node MUST also save the difference of `datagram_size` of the first fragment in the VRB, and add it to the `datagram_size` and to the `fragment_offset` of all the subsequent fragments for that datagram.

5. New Dispatch types and headers

This specification enables the 6LoWPAN fragmentation sublayer to provide an MTU up to 2048 bytes to the upper layer, which can be the 6LoWPAN Header Compression sublayer that is defined in the "Compression Format for IPv6 Datagrams" [RFC6282] specification. In order to achieve this, this specification enables the fragmentation and the reliable transmission of fragments over a multihop 6LoWPAN mesh network.

This specification provides a technique that is derived from MPLS in order to forward individual fragments across a 6LoWPAN route-over mesh. The datagram_tag is used as a label; it is locally unique to the node that is the source MAC address of the fragment, so together the MAC address and the label can identify the fragment globally. A node may build the datagram_tag in its own locally-significant way, as long as the selected tag stays unique to the particular datagram for the lifetime of that datagram. It results that the label does not need to be globally unique but also that it must be swapped at each hop as the source MAC address changes.

This specification extends RFC 4944 [RFC4944] with 4 new Dispatch types, for Recoverable Fragment (RFRAG) headers with or without Acknowledgment Request (RFRAG vs. RFRAG-ARQ), and for the RFRAG Acknowledgment back, with or without ECN Echo (RFRAG-ACK vs. RFRAG-ECHO).

(to be confirmed by IANA) The new 6LoWPAN Dispatch types use the Value Bit Pattern of 11 1010xx from page 0 [RFC8025], as follows:

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

Figure 1: Additional Dispatch Value Bit Patterns

In the following sections, the semantics of "datagram_tag" are unchanged from [RFC4944] Section 5.3. "Fragmentation Type and Header." and is compatible with the fragment forwarding operation described in [I-D.wattheyne-6lo-minimal-fragment].

5.1. Recoverable Fragment Dispatch type and Header

In this specification, the size and offset of the fragments are expressed on the compressed packet form as opposed to the uncompressed - native - packet form.

The first fragment is recognized by a sequence of 0; it carries its `fragment_size` and the `datagram_size` of the compressed packet, whereas the other fragments carry their `fragment_size` and `fragment_offset`. The last fragment for a datagram is recognized when its `fragment_offset` and its `fragment_size` add up to the `datagram_size`.

Recoverable Fragments are sequenced and a bitmap is used in the RFRAG Acknowledgment to indicate the received fragments by setting the individual bits that correspond to their sequence.




Figure 2: RFRAG Dispatch type and Header

E: 1 bit; Explicit Congestion Notification; the "E" flag is reset by the source of the fragment and set by intermediate routers to signal that this fragment experienced congestion along its path.

Fragment_size: 10 bit unsigned integer; the size of this fragment in a unit that depends on the MAC layer technology. For IEEE Std. 802.15.4, the unit is octet, and the maximum fragment size, which is constrained by the maximum frame size of 128 octet minus the overheads of the MAC and Fragment Headers, is not limited by this encoding.

X: 1 bit; Ack Requested: when set, the sender requires an RFRAG Acknowledgment from the receiver.

Sequence: 5 bit unsigned integer; the sequence number of the fragment. Fragments are sequence numbered [0..N] where N is in [0..31]. A sequence of 0 indicates the first fragment in a datagram. For IEEE Std. 802.15.4, as long as the overheads enable

a fragment size of 64 octets or more, this enables to fragment a packet of 2047 octets.

Fragment_offset: 16 bit unsigned integer;

- * When set to a non-0 value, the semantics of the Fragment_offset depends on the value of the Sequence.
 - + When the Sequence is not 0, this field indicates the offset of the fragment in the compressed form. The fragment should be forwarded based on an existing VRB as described in Section 7.2, or silently dropped if none is found.
 - + For a first fragment (i.e. with a sequence of 0), this field is overloaded to indicate the total_size of the compressed packet, to help the receiver allocate an adapted buffer for the reception and reassembly operations. This format limits the maximum MTU on a 6LoWPAN link to 2047 bytes, but 1280 bytes is the recommended value to avoid issues with IPV6 Path MTU Discovery [RFC8201]. The fragment should be routed based on the destination IPv6 address, and an VRB state should be installed as described in Section 7.1.
- * When set to 0, this field indicates an abort condition and all state regarding the datagram should be cleaned up once the processing of the fragment is complete; the processing of the fragment depends on whether there is a VRB already established for this datagram, and the next hop is still reachable:
 - + if a VRB already exists and is not broken, the fragment is to be forwarded along the associated Label Switched Path (LSP) as described in Section 7.2, but regardless of the value of the Sequence field;
 - + else, if the Sequence is 0, then the fragment is to be routed as described in Section 7.1 but no state is conserved afterwards.

If the fragment cannot be forwarded or routed, then an abort RFRAG-ACK is sent back to the source.

5.2. RFRAG Acknowledgment Dispatch type and Header

This specification also defines a 4-octet RFRAG Acknowledgment bitmap that is used by the reassembling end point to confirm selectively the reception of individual 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:

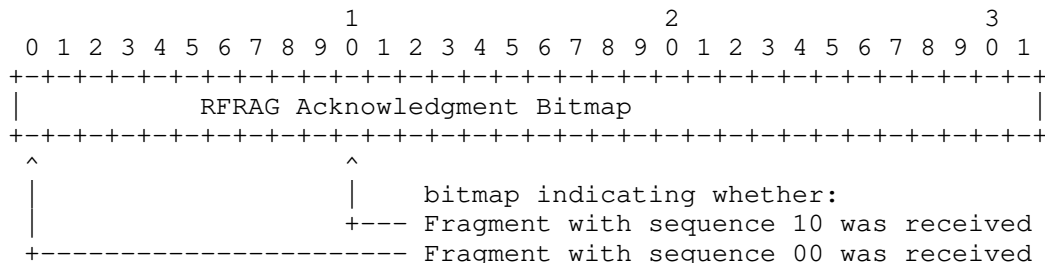


Figure 3: RFRAG Acknowledgment bitmap encoding

Figure 4 shows an example Acknowledgment bitmap which indicates that all fragments from sequence 0 to 20 were received, except for fragments 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 RFRAG Acknowledgment Bitmap is included in a RFRAG Acknowledgment header, as follows:




Figure 5: RFRAG Acknowledgment Dispatch type and Header

Y: 1 bit; Explicit Congestion Notification Echo

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.

RFRAG Acknowledgment Bitmap

An RFRAG Acknowledgment Bitmap, whereby setting the bit at offset *x* indicates that fragment *x* was received, as shown in Figure 3. All 0's is a NULL bitmap that indicates that the fragmentation process is aborted. All 1's is a FULL bitmap that indicates that the fragmentation process is complete, all fragments were received at the reassembly end point.

6. Fragments Recovery

The Recoverable Fragment headers RFRAG and RFRAG-ARQ are used to transport a fragment and optionally request an RFRAG Acknowledgment that will confirm the good reception of a one or more fragments. An RFRAG Acknowledgment can optionally carry an ECN indication; it is carried as a standalone header in a message that is sent back to the 6LoWPAN endpoint that was the source of the fragments, as known by its MAC address. The process ensures that at every hop, the source MAC address and the datagram_tag in the received fragment are enough information to send the RFRAG Acknowledgment back towards the source 6LoWPAN endpoint by reversing the MPLS operation.

The 6LoWPAN endpoint that fragments the packets at 6LoWPAN level (the sender) also controls when the reassembling end point sends the RFRAG Acknowledgments by setting the Ack Requested flag in the RFRAG packets. It may set the Ack Requested flag on any fragment to perform congestion control by limiting the number of outstanding fragments, which are the fragments that have been sent but for which reception or loss was not positively confirmed by the reassembling endpoint. When the sender of the fragment knows that an underlying link-layer mechanism protects the Fragments, it may refrain from using the RFRAG Acknowledgment mechanism, and never set the Ack Requested bit. When it receives a fragment with the ACK Request flag set, the 6LoWPAN endpoint that reassembles the packets at 6LoWPAN level (the receiver) sends back an RFRAG Acknowledgment to confirm reception of all the fragments it has received so far.

The sender transfers a controlled number of fragments and MAY flag the last fragment of a series with an RFRAG 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 MAY issue unsolicited acknowledgments. An unsolicited acknowledgment signals to the sender endpoint that it can resume sending if it had reached its maximum number of outstanding fragments. Another use is to inform that the reassembling endpoint has canceled 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.

An observation is that streamlining forwarding of fragments generally reduces the latency over the LLN mesh, providing room for retries within existing upper-layer reliability mechanisms. The sender protects the transmission over the LLN mesh with a retry timer that is computed according to the method detailed in [RFC6298]. It is expected that the upper layer retries obey the recommendations in "UDP Usage Guidelines" [RFC8085], 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 a single frequency is used by contiguous hops, the sender should wait a reasonable amount of time between fragments so as to let a fragment progress a few hops and avoid hidden terminal issues. This precaution is not required on channel hopping technologies such as Time Slotted CHannel Hopping (TSCH) [RFC6554]

When the sender decides that a packet should be dropped and the fragmentation process canceled, it sends a pseudo fragment with the `fragment_offset`, `sequence` and `fragment_size` all set to 0, 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 reassembly buffers, or considers that this packet is already fully reassembled and passed to the upper layer. In that case, the receiver SHOULD indicate so to the sender with a NULL bitmap in a RFRAG Acknowledgment. Upon an acknowledgment with a NULL bitmap, the

sender endpoint MUST abort the transmission of the fragmented datagram.

7. Forwarding Fragments

It is assumed that the first Fragment is large enough to carry the IPv6 header and make routing decisions. If that is not so, then this specification MUST NOT be used.

This specification extends the Virtual Reassembly Buffer (VRB) technique to forward fragments with no intermediate reconstruction of the entire packet. The first fragment carries the IP header and it is routed all the way from the fragmenting end point to the reassembling end point. Upon the first fragment, the routers along the path install a label-switched path (LSP), and the following 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. All fragments follow the same path and fragments are delivered in the order at which they are sent.

7.1. Upon the first fragment

In Route-Over mode, the source and destination MAC addressed in a frame change 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. Upon a first fragment (i.e. with a sequence of zero), a VRB and the associated LSP state are created for the tuple (source MAC address, `datagram_tag`) and the fragment is forwarded along the IPv6 route that matches the destination IPv6 address in the IPv6 header as prescribed by [I-D.watteyne-6lo-minimal-fragment]. The LSP state enables to match the (previous MAC address, `datagram_tag`) in an incoming fragment to the tuple (next MAC address, swapped `datagram_tag`) used in the forwarded fragment and points at the VRB. In addition, the router also forms a Reverse LSP state indexed by the MAC address of the next hop and the swapped `datagram_tag`. This reverse LSP state also points at the VRB and enables to match the (next MAC address, swapped `datagram_tag`) found in an RFRAG Acknowledgment to the tuple (previous MAC address, `datagram_tag`) used when forwarding a Fragment Acknowledgment (RFRAG-ACK) back to the sender endpoint.

7.2. Upon the next fragments

Upon a next fragment (i.e. with a non-zero sequence), the router looks up a LSP indexed by the tuple (MAC address, `datagram_tag`) found in the fragment. If it is found, the router forwards the fragment

using the associated VRB as prescribed by [I-D.watteyne-6lo-minimal-fragment].

if the VRB for the tuple is not found, the router builds an RFRAG-ACK to abort the transmission of the packet. The resulting message has the following information:

- o The source and destination MAC addresses are swapped from those found in the fragment
- o The datagram_tag set to the datagram_tag found in the fragment
- o A null bitmap is used to signal the abort condition

At this point the router is all set and can send the RFRAG-ACK back to the previous router. The RFRAG-ACK should normally be forwarded all the way to the source using the reverse LSP state in the VRBs in the intermediate routers as described in the next section.

7.3. Upon the RFRAG Acknowledgments

Upon an RFRAG-ACK, the router looks up a Reverse LSP indexed by the tuple (MAC address, datagram_tag), which are respectively the source MAC address of the received frame and the received datagram_tag. If it is found, the router forwards the fragment using the associated VRB as prescribed by [I-D.watteyne-6lo-minimal-fragment], but using the Reverse LSP so that the RFRAG-ACK flows back to the sender endpoint.

If the Reverse LSP is not found, the router MUST silently drop the RFRAG-ACK message.

Either way, if the RFRAG-ACK indicates either an error (NULL bitmap) or that the fragment was entirely received (FULL bitmap), arms a short timer, and upon timeout, the VRB and all associate state are destroyed. During that time, fragments of that datagram may still be received, e.g. if the RFRAG-ACK was lost on the way back and the source retried the last fragment. In that case, the router sends an abort RFRAG-ACK along the Reverse LSP to complete the clean up.

8. 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].

9. IANA Considerations

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

10. Acknowledgments

The author wishes to thank Thomas Watteyne and Michael Richardson for in-depth reviews and comments. Also many thanks to Jonathan Hui, Jay Werb, Christos Polyzois, Soumitri Kolavennu, Pat Kinney, Margaret Wasserman, Richard Kelsey, Carsten Bormann and Harry Courtice for their various contributions.

11. References

11.1. Normative References

- [I-D.watteyne-6lo-minimal-fragment]
Watteyne, T., Bormann, C., and P. Thubert, "LLN Minimal Fragment Forwarding", draft-watteyne-6lo-minimal-fragment-02 (work in progress), July 2018.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [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>>.
- [RFC6554] Hui, J., Vasseur, JP., Culler, D., and V. Manral, "An IPv6 Routing Header for Source Routes with the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6554, DOI 10.17487/RFC6554, March 2012, <<https://www.rfc-editor.org/info/rfc6554>>.
- [RFC8025] Thubert, P., Ed. and R. Cragie, "IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Paging Dispatch", RFC 8025, DOI 10.17487/RFC8025, November 2016, <<https://www.rfc-editor.org/info/rfc8025>>.

- [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>>.
- [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>>.

11.2. 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-19 (work in progress), December 2018.
- [IEEE.802.15.4] IEEE, "IEEE Standard for Low-Rate Wireless Networks", IEEE Standard 802.15.4, DOI 10.1109/IEEE P802.15.4-REVd/D01, <<http://ieeexplore.ieee.org/document/7460875/>>.
- [RFC2914] Floyd, S., "Congestion Control Principles", BCP 41, RFC 2914, DOI 10.17487/RFC2914, September 2000, <<https://www.rfc-editor.org/info/rfc2914>>.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, DOI 10.17487/RFC3031, January 2001, <<https://www.rfc-editor.org/info/rfc3031>>.
- [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, <<https://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, <<https://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, <<https://www.rfc-editor.org/info/rfc4963>>.

- [RFC5681] Allman, M., Paxson, V., and E. Blanton, "TCP Congestion Control", RFC 5681, DOI 10.17487/RFC5681, September 2009, <<https://www.rfc-editor.org/info/rfc5681>>.
- [RFC6298] Paxson, V., Allman, M., Chu, J., and M. Sargent, "Computing TCP's Retransmission Timer", RFC 6298, DOI 10.17487/RFC6298, June 2011, <<https://www.rfc-editor.org/info/rfc6298>>.
- [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>>.
- [RFC7554] Watteyne, T., Ed., Palattella, M., and L. Grieco, "Using IEEE 802.15.4e Time-Slotted Channel Hopping (TSCH) in the Internet of Things (IoT): Problem Statement", RFC 7554, DOI 10.17487/RFC7554, May 2015, <<https://www.rfc-editor.org/info/rfc7554>>.
- [RFC7567] Baker, F., Ed. and G. Fairhurst, Ed., "IETF Recommendations Regarding Active Queue Management", BCP 197, RFC 7567, DOI 10.17487/RFC7567, July 2015, <<https://www.rfc-editor.org/info/rfc7567>>.
- [RFC8085] Eggert, L., Fairhurst, G., and G. Shepherd, "UDP Usage Guidelines", BCP 145, RFC 8085, DOI 10.17487/RFC8085, March 2017, <<https://www.rfc-editor.org/info/rfc8085>>.
- [RFC8087] Fairhurst, G. and M. Welzl, "The Benefits of Using Explicit Congestion Notification (ECN)", RFC 8087, DOI 10.17487/RFC8087, March 2017, <<https://www.rfc-editor.org/info/rfc8087>>.
- [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>>.

Appendix A. 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:

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.

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.

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, the lack of recovery in the original fragmentation system of RFC 4944 implies that 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.

Considering that RFC 4944 defines an MTU is 1280 bytes and that in most incarnations (but 802.15.4g) a IEEE Std. 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.

Mechanisms such as TCP or application-layer segmentation could 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. In addition, deploying such a mechanism requires that the end-to-end transport is aware of the delivery properties of the underlying LLN, which is a layer violation, and difficult to achieve from the far end of the IPv6 network.

Appendix B. Requirements

For one-hop communications, a number of Low Power and Lossy Network (LLN) link-layers propose a local acknowledgment mechanism that is enough to detect and recover the loss of 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 that may be multiple hops away. The method addresses the following requirements of a LLN:

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 new end-to-end fragment 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.

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.

Appendix C. Considerations On Flow Control

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 5.2.

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 [RFC7554] (TSCH) mode of operation of IEEE802.14.5, a fragment is forwarded over a different channel at a different time and it makes full sense to transmit the next fragment as soon as the previous fragment has had its chance to be forwarded at the next hop.

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.

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 [RFC7567] 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 reducing the number of outstanding fragments over a congested path by throttling the sources.

Section 6 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.

Author's Address

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

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

6lo
Internet-Draft
Intended status: Informational
Expires: September 12, 2019

T. Watteyne, Ed.
Analog Devices
C. Bormann
Universitaet Bremen TZI
P. Thubert
Cisco
March 11, 2019

LLN Minimal Fragment Forwarding
draft-ietf-6lo-minimal-fragment-01

Abstract

This document gives an overview of LLN Minimal Fragment Forwarding. When employing adaptation layer fragmentation in 6LoWPAN, it may be beneficial for a forwarder not to have to reassemble each packet in its entirety before forwarding it. This has always been possible with the original fragmentation design of RFC4944. This document is a companion document to [I-D.ietf-lwig-6lowpan-virtual-reassembly], which details the virtual Reassembly Buffer (VRB) implementation technique which reduces the latency and increases end-to-end reliability in route-over forwarding.

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 September 12, 2019.

Copyright Notice

Copyright (c) 2019 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. Overview of 6LoWPAN Fragmentation	2
2. Limits of Per-Hop Fragmentation and Reassembly	4
2.1. Latency	4
2.2. Memory Management and Reliability	4
3. Virtual Reassembly Buffer (VRB) Implementation	5
4. Security Considerations	5
5. IANA Considerations	6
6. Acknowledgments	6
7. Informative References	6
Authors' Addresses	6

1. Overview of 6LoWPAN Fragmentation

6LoWPAN fragmentation is defined in [RFC4944]. Although [RFC6282] updates [RFC4944], it does not redefine 6LoWPAN fragmentation.

We use Figure 1 to illustrate 6LoWPAN fragmentation. We assume node A forwards a packet to node B, possibly as part of a multi-hop route between IPv6 source and destination nodes which are neither A nor B.




Figure 1: Fragmentation at node A, reassembly at node B.

Node A starts by compacting the IPv6 packet using the header compression mechanism defined in [RFC6282]. If the resulting 6LoWPAN packet does not fit into a single link-layer frame, node A's 6LoWPAN

sublayer cuts it into multiple 6LoWPAN fragments, which it transmits as separate link-layer frames to node B. Node B's 6LoWPAN sublayer reassembles these fragments, inflates the compressed header fields back to the original IPv6 header, and hands over the full IPv6 packet to its IPv6 layer.

In Figure 1, a packet forwarded by node A to node B is cut into nine fragments, numbered 1 to 9. Each fragment is represented by the '#' symbol. Node A has sent fragments 1, 2, 3, 5, 6 to node B. Node B has received fragments 1, 2, 3, 6 from node A. Fragment 5 is still being transmitted at the link layer from node A to node B.

Conceptually, a reassembly buffer for 6LoWPAN contains:

- o a `datagram_size`,
- o a `datagram_tag`, associated to the link-layer sender and receiver addresses to which the `datagram_tag` is local,
- o the actual packet data from the fragments received so far, in a form that makes it possible to detect when the whole packet has been received and can be processed or forwarded,
- o a timer that allows discarding a partially reassembled packet after some timeout.

A fragmentation header is added to each fragment; it indicates what portion of the packet that fragment corresponds to. Section 5.3 of [RFC4944] defines the format of the header for the first and subsequent fragments. All fragments are tagged with a 16-bit "datagram_tag", used to identify which packet each fragment belongs to. Each fragment can be uniquely identified by the source and destination link-layer addresses of the frame that carries it, and the `datagram_tag`. The value of the `datagram_tag` only needs to be locally unique to nodes A and B.

Node B's typical behavior, per [RFC4944], is as follows. Upon receiving a fragment from node A with a `datagram_tag` previously unseen from node A, node B allocates a buffer large enough to hold the entire packet. The length of the packet is indicated in each fragment (the `datagram_size` field), so node B can allocate the buffer even if the first fragment it receives is not fragment 1. As fragments come in, node B fills the buffer. When all fragments have been received, node B inflates the compressed header fields into an IPv6 header, and hands the resulting IPv6 packet to the IPv6 layer.

This behavior typically results in per-hop fragmentation and reassembly. That is, the packet is fully reassembled, then (re)fragmented, at every hop.

2. Limits of Per-Hop Fragmentation and Reassembly

There are at least 2 limits to doing per-hop fragmentation and reassembly. See [ARTICLE] for detailed simulation results on both limits.

2.1. Latency

When reassembling, a node needs to wait for all the fragments to be received before being able to generate the IPv6 packet, and possibly forward it to the next hop. This repeats at every hop.

This may result in increased end-to-end latency compared to a case where each fragment is forwarded without per-hop reassembly.

2.2. Memory Management and Reliability

Constrained nodes have limited memory. Assuming 1 kB reassembly buffers, typical nodes only have enough memory for 1-3 reassembly buffers.

Assuming the topology from Figure 2, where nodes A, B, C and D all send packets through node E. We further assume that node E's memory can only hold 3 reassembly buffers.




Figure 2: Illustrating the Memory Management Issue.

When nodes A, B and C concurrently send fragmented packets, all 3 reassembly buffers in node E are occupied. If, at that moment, node D also sends a fragmented packet, node E has no option but to drop one of the packets, lowering end-to-end reliability.

3. Virtual Reassembly Buffer (VRB) Implementation

Virtual Reassembly Buffer (VRB) is the implementation technique described in [I-D.ietf-lwig-6lowpan-virtual-reassembly] in which a forwarder does not reassemble each packet in its entirety before forwarding it.

VRB overcomes the limits listed in Section 2. Nodes don't wait for the last fragment before forwarding, reducing end-to-end latency. Similarly, the memory footprint of VRB is just the VRB table, reducing the packet drop probability significantly.

There are, however, limits:

- Non-zero Packet Drop Probability: Each VRB table entry can be 12 B (assuming 16-bit link-layer addresses). This is a footprint 2 orders of magnitude smaller compared to needing a 1280-byte reassembly buffer for each packet. Yet, the size of the VRB table necessarily remains finite. In the extreme case where a node is required to concurrently forward more packets than it has entries in its VRB table, packets are dropped.
- No Fragment Recovery: There is no mechanism in VRB for the node that reassembles a packet to request a single missing fragment. Dropping a fragment requires the whole packet to be resent. This causes unnecessary traffic, as fragments are forwarded even when the destination node can never construct the original IPv6 packet.
- No Per-Fragment Routing: All subsequent fragments follow the same sequence of hops from the source to the destination node as fragment 1.

The severity and occurrence of these limits depends on the link-layer used. Whether these limits are acceptable depends entirely on the requirements the application places on the network.

If the limits are present and not acceptable for the application, future specifications may define new protocols to overcome these limits. One example is [I-D.thubert-6lo-fragment-recovery] which defines a protocol which allows fragment recovery.

4. Security Considerations

An attacker can perform a DoS attack on a node implementing VRB by generating a large number of bogus "fragment 1" fragments without sending subsequent fragments. This causes the VRB table to fill up.

Secure joining and the link-layer security that it sets up protects against those attacks from network outsiders.

5. IANA Considerations

No requests to IANA are made by this document.

6. Acknowledgments

The authors would like to thank Yasuyuki Tanaka for his in-depth review of this document.

7. Informative References

- [ARTICLE] Tanaka, Y., Minet, P., and T. Watteyne, "6LoWPAN Fragment Forwarding", IEEE Communications Standards Magazine , 2009.
- [BOOK] Shelby, Z. and C. Bormann, "6LoWPAN", John Wiley & Sons, Ltd monograph, DOI 10.1002/9780470686218, November 2009.
- [I-D.ietf-lwig-6lowpan-virtual-reassembly]
Bormann, C. and T. Watteyne, "Virtual reassembly buffers in 6LoWPAN", draft-ietf-lwig-6lowpan-virtual-reassembly-00 (work in progress), July 2018.
- [I-D.thubert-6lo-fragment-recovery]
Thubert, P., "6LoWPAN Selective Fragment Recovery", draft-thubert-6lo-fragment-recovery-01 (work in progress), June 2018.
- [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>>.

Authors' Addresses

Thomas Watteyne (editor)
Analog Devices
32990 Alvarado-Niles Road, Suite 910
Union City, CA 94587
USA

Email: thomas.watteyne@analog.com

Carsten Bormann
Universitaet Bremen TZI
Postfach 330440
Bremen D-28359
Germany

Email: cabo@tzi.org

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

Email: pthubert@cisco.com

6Lo Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 5, 2019

J. Hou
B. Liu
Huawei Technologies
Y-G. Hong
ETRI
X. Tang
SGEPRI
C. Perkins
Futurewei
February 1, 2019

Transmission of IPv6 Packets over PLC Networks
draft-ietf-6lo-plc-00

Abstract

Power Line Communication (PLC), namely using the electric-power lines for indoor and outdoor communications, has been widely applied to support Advanced Metering Infrastructure (AMI), especially smart meters for electricity. The inherent advantage of existing electricity infrastructure facilitates the expansion of PLC deployments, and moreover, a wide variety of accessible devices raises the potential demand of IPv6 for future applications. This document describes how IPv6 packets are transported over constrained PLC networks, such as ITU-T G.9903, IEEE 1901.1 and IEEE 1901.2.

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 August 5, 2019.

Copyright Notice

Copyright (c) 2019 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. Requirements Notation and Terminology	3
3. Overview of PLC	5
3.1. Protocol Stack	5
3.2. Addressing Modes	6
3.3. Maximum Transmission Unit	6
3.4. Routing Protocol	7
4. IPv6 over PLC	7
4.1. Stateless Address Autoconfiguration	7
4.2. IPv6 Link Local Address	8
4.3. Unicast Address Mapping	9
4.3.1. Unicast Address Mapping for IEEE 1901.1	9
4.3.2. Unicast Address Mapping for IEEE 1901.2 and ITU-T G.9903	10
4.4. Neighbor Discovery	10
4.5. Header Compression	11
4.6. Fragmentation and Reassembly	11
4.7. Extension at 6lo Adaptation Layer	12
5. Internet Connectivity Scenarios and Topologies	13
6. IANA Considerations	16
7. Security Consideration	16
8. Acknowledgements	16
9. References	16
9.1. Normative References	16
9.2. Informative References	18
Authors' Addresses	19

1. Introduction

The idea of using power lines for both electricity supply and communication can be traced back to the beginning of the last century. With the advantage of existing power grid, Power Line Communication (PLC) is a good candidate for supporting various service scenarios such as in houses and offices, in trains and vehicles, in smart grid and advanced metering infrastructure (AMI). The data acquisition devices in these scenarios share common features such as fixed position, large quantity, low data rate and low power consumption.

Although PLC technology has evolved over several decades, it has not been fully adapted for IPv6 based constrained networks. The 6Lo related scenarios lie in the low voltage PLC networks with most applications in the area of Advanced Metering Infrastructure (AMI), Vehicle-to-Grid communications, in-home energy management and smart street lighting. IPv6 is important for PLC networks, due to its large address space and efficient address auto-configuration. A comparison among various existing PLC standards is provided to facilitate the selection of the most applicable standard in particular scenarios.

This specification provides a brief overview of PLC technologies. Some of them have LLN characteristics, i.e. limited power consumption, memory and processing resources. This specification is focused on the transmission of IPv6 packets over those "constrained" PLC networks. The general approach is to adapt elements of the 6LoWPAN specifications [RFC4944], [RFC6282], and [RFC6775] to constrained PLC networks. Compared to [I-D.popa-6lo-6loplc-ipv6-over-ieee19012-networks], this document provides a structured and greatly expanded specification of an adaptation layer for IPv6 over PLC (6LoPLC) networks.

2. Requirements Notation and Terminology

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

This document often uses the following acronyms and terminologies:

6LoWPAN: IPv6 over Low-Power Wireless Personal Area Network

AMI: Advanced Metering Infrastructure

BBPLC: Broadband Power Line Communication

CID: Context ID

Coordinator: A device capable of relaying messages.

DAD: Duplicate Address Detection

PAN device: An entity follows the PLC standards and implements the protocol stack described in this draft.

EV: Electric Vehicle

IID: IPv6 Interface Identifier

IPHC: IP Header Compression

LAN: Local Area Network

MSDU: MAC Service Data Unit

MTU: Maximum Transmission Unit

NBPLC: Narrowband Power Line Communication

OFDM: Orthogonal Frequency Division Multiplexing

PANC: PAN Coordinator, a coordinator which also acts as the primary controller of a PAN.

PLC: Power Line Communication

PSDU: PHY Service Data Unit

RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks

RA: Router Advertisement

WAN: Wide Area Network

The terminology used in this draft is aligned with IEEE 1901.2

IEEE 1901.2	IEEE 1901.1	ITU-T G.9903
PAN Coordinator	Central Coordinator	PAN Coordinator
Coordinator	Proxy Coordinator	Full-function device
Device	Station	PAN Device

Table 1: Terminology Mapping between PLC standards

3. Overview of PLC

PLC technology enables convenient two-way communications for home users and utility companies to monitor and control electric plugged devices such as electricity meters and street lights. Due to the large range of communication frequencies, PLC is generally classified into two categories: Narrowband PLC (NBPLC) for automation of sensors (which have low frequency band and low power cost), and Broadband PLC (BBPLC) for home and industry networking applications. Various standards have been addressed on the MAC and PHY layers for this communication technology, e.g. BBPLC (1.8–250 MHz) including IEEE 1901 and ITU-T G.hn, and NBPLC (3–500 kHz) including ITU-T G.9902 (G.hnem), ITU-T G.9903 (G3-PLC) [ITU-T_G.9903], ITU-T G.9904 (PRIME), IEEE 1901.2 [IEEE_1901.2] (combination of G3-PLC and PRIME PLC) and IEEE 1901.2a [IEEE_1901.2a] (an amendment to IEEE 1901.2). Moreover, recently a new PLC standard IEEE 1901.1 [IEEE_1901.1], which aims at the medium frequency band less than 12 MHz, has been published by the IEEE standard for Smart Grid Powerline Communication Working Group (SGPLC WG). IEEE 1901.1 balances the needs for bandwidth versus communication range, and is thus a promising option for 6Lo applications. Currently, this specification is focused on IEEE 1901.1, IEEE 1901.2 and ITU-T G.9903.

3.1. Protocol Stack

The protocol stack for IPv6 over PLC is illustrated in Figure 1. The PLC MAC/PHY layer corresponds to IEEE 1901.1, IEEE 1901.2 or ITU-T G.9903. The 6Lo adaptation layer for PLC is illustrated in Section 4. For multihop tree and mesh topologies, a routing protocol is likely to be necessary. The routes can be built in mesh-under mode at layer 2 or in route-over mode at layer 3.




Figure 1: PLC Protocol Stack

3.2. Addressing Modes

Each PLC device has a globally unique long address of 48-bit ([IEEE_1901.1]) or 64-bit ([IEEE_1901.2], [ITU-T_G.9903]) and a short address of 12-bit ([IEEE_1901.1]) or 16-bit ([IEEE_1901.2], [ITU-T_G.9903]). The long address is set by the manufacturer according to the IEEE EUI-48 MAC address or the IEEE EUI-64 address. Each PLC device joins the network by using the long address and communicates with other devices by using the short address after joining the network.

3.3. Maximum Transmission Unit

The Maximum Transmission Unit (MTU) of the MAC layer determines whether fragmentation and reassembly are needed at the adaptation layer of IPv6 over PLC. IPv6 requires an MTU of 1280 octets or greater; thus for a MAC layer with MTU lower than this limit, fragmentation and reassembly at the adaptation layer are required.

The IEEE 1901.1 MAC supports upper layer packets up to 2031 octets. The IEEE 1901.2 MAC layer supports the MTU of 1576 octets (the original value of 1280 bytes was updated in 2015 [IEEE_1901.2a]). Though fragmentation and reassembly are not needed in these two technologies, other 6lo functions like header compression are still applicable and useful, particularly in high-noise communication environments.

The MTU for ITU-T G.9903 is 400 octets, insufficient for supporting IPv6's MTU. For this reason, fragmentation and reassembly as per [RFC4944] MUST be enabled for G.9903-based networks.

3.4. Routing Protocol

Routing protocols suitable for use in PLC networks include:

- o RPL (Routing Protocol for Low-Power and Lossy Networks) [RFC6550] is a layer 3 routing protocol. AODV-RPL [I-D.ietf-roll-aodv-rpl] updates RPL to include reactive, point-to-point, and asymmetric routing. IEEE 1901.2 specifies Information Elements (IEs) with MAC layer metrics, which can be provided to L3 routing protocol for parent selection. For IPv6-addressable PLC networks, a layer-3 routing protocol such as RPL and/or AODV-RPL SHOULD be supported in the standard.
- o IEEE 1901.1 supports L2 routing. Each PLC node maintains a L2 routing table, in which each route entry comprises the short addresses of the destination and the related next hop. The route entries are built during the network establishment via a pair of association request/confirmation messages. The route entries can be changed via a pair of proxy change request/confirmation messages. These association and proxy change messages MUST be approved by the central coordinator.
- o LOADng is a reactive protocol operating at layer 2 or layer 3. Currently, LOADng is supported in ITU-T G.9903 [ITU-T_G.9903], and the IEEE 1901.2 standard refers to ITU-T G.9903 for LOAD-based networks.

4. IPv6 over PLC

6LoWPAN standards [RFC4944], [RFC6775], and [RFC6282] provides useful functionality including link-local IPv6 addresses, stateless address auto-configuration, neighbor discovery and header compression. However, due to the different characteristics of the PLC media, the 6LoWPAN adaptation layer cannot perfectly fulfill the requirements. Besides, some of the features like fragmentation and reassembly are redundant to some PLC technologies. These considerations suggest the need for a dedicated adaptation layer for PLC, which is detailed in the following subsections.

4.1. Stateless Address Autoconfiguration

To obtain an IPv6 Interface Identifier (IID), a PLC device performs stateless address autoconfiguration [RFC4944]. The autoconfiguration can be based on either a long or short link-layer address.

The IID can be based on the device's 48-bit MAC address or its EUI-64 identifier [EUI-64]. A 48-bit MAC address MUST first be extended to a 64-bit Interface ID by inserting 0xFFFE at the fourth and fifth

octets as specified in [RFC2464]. The IPv6 IID is derived from the 64-bit Interface ID by inverting the U/L bit [RFC4291].

For IEEE 1901.2 and ITU-T G.9903, a 48-bit "pseudo-address" is formed by the 16-bit PAN ID, 16 zero bits and the 16-bit short address. Then, the 64-bit Interface ID MUST be derived by inserting 16-bit 0xFFFE into as follows:

```
16_bit_PAN:00FF:FE00:16_bit_short_address
```

For the 12-bit short addresses used by IEEE 1901.1, the 48-bit pseudo-address is formed by 24-bit NID (Network Identifier, YYYYYY), 12 zero bits and a 12-bit TEI (Terminal Equipment Identifier, XXX). The 64-bit Interface ID MUST be derived by inserting 16-bit 0xFFFE into this 48-bit pseudo-address as follows:

```
YYYY:YYFF:FE00:0XXX
```

Since the derived Interface ID is not global, the "Universal/Local" (U/L) bit (7th bit) and the Individual/Group bit (8th bit) MUST both be set to zero. In order to avoid any ambiguity in the derived Interface ID, these two bits MUST NOT be used to generate the PANID (for IEEE 1901.2 and ITU-T G.9903) or NID (for IEEE 1901.1). In other words, the PANID or NID MUST always be chosen so that these bits are zeros.

For privacy reasons, the IID derived by the MAC address SHOULD only be used for link-local address configuration. A PLC host SHOULD use the IID derived by the link-layer short address to configure the IPv6 address used for communication with the public network; otherwise, the host's MAC address is exposed.

4.2. IPv6 Link Local Address

The IPv6 link-local address [RFC4291] for a PLC interface is formed by appending the IID, as defined above, to the prefix FE80::/64 (see Figure 2).




Figure 2: IPv6 Link Local Address for a PLC interface

4.3. Unicast Address Mapping

The address resolution procedure for mapping IPv6 unicast addresses into PLC link-layer addresses follows the general description in section 7.2 of [RFC4861]. [RFC6775] improves this procedure by eliminating usage of multicast NS. The resolution is realized by the NCEs (neighbor cache entry) created during the address registration at the routers. 6775-update further improves the registration procedure by enabling multiple LLNs to form an IPv6 subnet, and by inserting a link-local address registration to better serve proxy registration of new devices.

4.3.1. Unicast Address Mapping for IEEE 1901.1

The Source/Target Link-layer Address options for IEEE_1901.1 used in the Neighbor Solicitation and Neighbor Advertisement have the following form.




Figure 3: Unicast Address Mapping for IEEE 1901.1

Option fields:

Type: 1 for Source Link-layer Address and 2 for Target Link-layer Address.

Length: The length of this option (including type and length fields) in units of 8 octets. The value of this field is 1 for the 12-bit IEEE 1901.1 PLC short addresses.

NID: 24-bit Network IDentifier

Padding: 12 zero bits

TEI: 12-bit Terminal Equipment Identifier

In order to avoid the possibility of duplicated IPv6 addresses, the value of the NID MUST be chosen so that the 7th and 8th bits of the first byte of the NID are both zero.

4.3.2. Unicast Address Mapping for IEEE 1901.2 and ITU-T G.9903

The Source/Target Link-layer Address options for IEEE_1901.2 and ITU-T G.9903 used in the Neighbor Solicitation and Neighbor Advertisement have the following form.




Figure 4: Unicast Address Mapping for IEEE 1901.2

Option fields:

Type: 1 for Source Link-layer Address and 2 for Target Link-layer Address.

Length: The length of this option (including type and length fields) in units of 8 octets. The value of this field is 1 for the 16-bit IEEE 1901.2 PLC short addresses.

PAN ID: 16-bit PAN Identifier

Padding: 16 zero bits

Short Address: 16-bit short address

In order to avoid the possibility of duplicated IPv6 addresses, the value of the PAN ID MUST be chosen so that the 7th and 8th bits of the first byte of the PAN ID are both zero.

4.4. Neighbor Discovery

Neighbor discovery procedures for 6LoWPAN networks are described in Neighbor Discovery Optimization for 6LoWPANs [RFC6775] and [I-D.ietf-6lo-rfc6775-update]. These optimizations support the registration of sleeping hosts. Although PLC devices are electrically powered, sleeping mode SHOULD still be used for power saving.

For IPv6 address prefix dissemination, Router Solicitations (RS) and Router Advertisements (RA) MAY be used as per [RFC6775]. If the PLC network uses route-over mesh, the IPv6 prefix MAY be disseminated by the layer 3 routing protocol, such as RPL which includes the prefix

in the DIO message. In this case, the prefix information option (PIO) MUST NOT be included in the Router Advertisement.

For context information dissemination, Router Advertisements (RA) MUST be used as per [RFC6775]. The 6LoWPAN context option (6CO) MUST be included in the RA to disseminate the Context IDs used for prefix compression.

For address registration, a PLC host MUST register its address to the router using Neighbor Solicitation and Neighbor Advertisement messages. RFC6775-update PLC devices MUST include the EARO with the 'R' flag set when sending Neighbor Solicitations, and process Neighbor Advertisements that include EARO to extract status information. If DHCPv6 is used to assign addresses, or the IPv6 address is derived by unique long or short link layer address, Duplicate Address Detection (DAD) MUST NOT be utilized. Otherwise, DAD MUST be performed: RFC6775-only PLC devices MUST perform multihop DAD against a 6LBR by using DAR and DAC messages, while for RFC6775-update devices, DAD is proxied by a routing registrar, which MAY operate according to Optimistic DAD (ODAD) [RFC4429].

The mesh-under ITU-T G.9903 network SHOULD NOT utilize the address registration as described in [RFC6775]. ITU-T G.9903 PLC networks MUST use the 6LoWPAN Context Option (6CO) specified in [RFC6775] (see clause 9.4.1.1 in [ITU-T_G.9903]), which can be attached in Router Advertisements to disseminate Context IDs (CIDs) to use for compressing prefixes.

4.5. Header Compression

The compression of IPv6 datagrams within PLC MAC frames refers to [RFC6282], which updates [RFC4944]. Header compression as defined in [RFC6282] which specifies the compression format for IPv6 datagrams on top of IEEE 802.15.4, is included in this document as the basis for IPv6 header compression in PLC. For situations when PLC MAC MTU cannot support the 1280-octet IPv6 packet, headers MUST be compressed according to [RFC6282] encoding formats.

4.6. Fragmentation and Reassembly

PLC differs from other wired technologies in that the communication medium is not shielded; thus, to successfully transmit data through power lines, PLC Data Link layer provides the function of segmentation and reassembly. A Segment Control Field is defined in the MAC frame header regardless of whether segmentation is required. The number of data octets of the PHY payload can change dynamically based on channel conditions, thus the MAC payload segmentation in the MAC sublayer is enabled and guarantees a reliable one-hop data

transmission. Fragmentation and reassembly is still required at the adaptation layer, if the MAC layer cannot support the minimum MTU demanded by IPv6, which is 1280 octets.

In IEEE 1901.1 and IEEE 1901.2, since the MAC layer supports payloads of 2031 octets and 1576 octets respectively, fragmentation is not needed for IPv6 packet transmission. The fragmentation and reassembly defined in [RFC4944] SHOULD NOT be used in the 6lo adaptation layer of IEEE 1901.2.

In ITU-T G.9903, the maximum MAC payload size is fixed to 400 octets, so to cope with the required MTU of 1280 octets by IPv6, fragmentation and reassembly at 6lo adaptation layer MUST be provided referring to [RFC4944].

4.7. Extension at 6lo Adaptation Layer

Apart from the Dispatch and LOWPAN_IPHC headers specified in [RFC4944], an additional Command Frame Header is needed for the mesh routing procedure in LOADng protocol. Figure 5 illustrates the format of the Command Frame Header [RFC8066]. The ESC dispatch type (01000000b) indicates an ESC extension type follows (see [RFC4944] and [RFC6282]). Then this 1-octet dispatch field is used as the Command Frame Header and filled with the Command ID. The Command ID can be classified into 4 types:

- o LOADng message (0x01)
- o LoWPAN bootstrapping protocol message (0x02)
- o Reserved by ITU-T (0x03-0x0F)
- o CMSR protocol messages (0x10-0x1F)

The LOADng message is used to provide the default routing protocol LOADng while the LoWPAN bootstrapping protocol message is for the LoWPAN bootstrap procedure. The CMSR protocol messages are specified for the Centralized metric-based source routing [ITU-T G.9905] which is out of the scope of this draft.




Figure 5: Command Frame Header Format of ITU-T G.9903

Command Frame Header appears in the last position if more than one header is present in the 6LoWPAN frame [ITU-T_G.9903]. On the other hand, this Command Frame Header MUST appear before the LoWPAN_IPHC dispatch type as per[RFC8066].

- o Regarding the order of the command frame header, the inconsistency between G.9903 and RFC8066 still exists and is being solved in ITU-T SG15/Q15.

Following these two requirements of header order mentioned above, an example of the header order is illustrated in Figure 6 including the Fragmentation type, Fragmentation header, ESC dispatch type, ESC Extension Type (Command ID), ESC Dispatch Payload (Command Payload), LoWPAN_IPHC Dispatch Type, LoWPAN_IPHC header, and Payload.

```
+-----+-----+-----+-----+-----+-----+-----+-----+
|F typ|F hdr| ESC | EET | EDP |Dispatch|LoWPAN_IPHC hdr| Payld|
+-----+-----+-----+-----+-----+-----+-----+-----+
```

Figure 6: A 6LoWPAN packet including the Command Frame Header

5. Internet Connectivity Scenarios and Topologies

The network model can be simplified to two kinds of network devices: PAN Coordinator (PANC) and PAN Device. The PANC is the primary coordinator of the PLC subnet and can be seen as a master node; PAN Devices are typically PLC meters and sensors. The PANC also serves as the Routing Registrar for proxy registration and DAD procedures, making use of the updated registration procedures in [I-D.ietf-6lo-rfc6775-update]. IPv6 over PLC networks are built as tree, mesh or star according to the use cases. Every network requires at least one PANC to communicate with each PAN Device. Note that the PLC topologies in this section are based on logical connectivity, not physical links.

The star topology is common in current PLC scenarios. In single-hop star topologies, communication at the link layer only takes place between a PAN Device and a PANC. The PANC typically collects data (e.g. a meter reading) from the PAN devices, and then concentrates and uploads the data through Ethernet or LPWAN (see Figure 7). The collected data is transmitted by the smart meters through PLC, aggregated by a concentrator, sent to the utility and then to a Meter Data Management System for data storage, analysis and billing. This topology has been widely applied in the deployment of smart meters, especially in apartment buildings.




Figure 7: PLC Star Network connected to the Internet

A tree topology is useful when the distance between a device A and PANC is beyond the PLC allowed limit and there is another device B in between able to communicate with both sides. Device B in this case acts both as a PAN Device and a Coordinator. For this scenario, the link layer communications take place between device A and device B, and between device B and PANC. An example of PLC tree network is depicted in Figure 8. This topology can be applied in the smart street lighting, where the lights adjust the brightness to reduce energy consumption while sensors are deployed on the street lights to provide information such as light intensity, temperature, humidity. Data transmission distance in the street lighting scenario is normally above several kilometers thus the PLC tree network is required. A more sophisticated AMI network may also be constructed into the tree topology which is depicted in [RFC8036]. A tree topology is suitable for AMI scenarios that require large coverage but low density, e.g. the deployment of smart meters in rural areas. RPL is suitable for maintenance of a tree topology in which there is no need for communication directly between PAN devices.




Figure 8: PLC Tree Network connected to the Internet

Mesh networking in PLC is of great potential applications and has been studied for several years. By connecting all nodes with their neighbors in communication range (see Figure 9), mesh topology dramatically enhances the communication efficiency and thus expands the size of PLC networks. A simple use case is the smart home scenario where the ON/OFF state of air conditioning is controlled by the state of home lights (ON/OFF) and doors (OPEN/CLOSE). AODV-RPL enables direct PAN device to PAN device communication, without being obliged to transmit frames through the PANC, which is a requirement often cited for AMI infrastructure.




Figure 9: PLC Mesh Network connected to the Internet

6. IANA Considerations

There are no IANA considerations related to this document.

7. Security Consideration

Due to the high accessibility of power grid, PLC might be susceptible to eavesdropping within its communication coverage, e.g. one apartment tenant may have the chance to monitor the other smart meters in the same apartment building. For security consideration, link layer security is guaranteed in every PLC technology.

IP addresses may be used to track devices on the Internet; such devices can in turn be linked to individuals and their activities. Depending on the application and the actual use pattern, this may be undesirable. To impede tracking, globally unique and non-changing characteristics of IP addresses should be avoided, e.g., by frequently changing the global prefix and avoiding unique link-layer derived IIDs in addresses. [RFC3315], [RFC3972], [RFC4941], [RFC5535], [RFC7217], and [RFC8065] provide valuable information for IID formation with improved privacy, and are RECOMMENDED for IPv6 networks.

8. Acknowledgements

We gratefully acknowledge suggestions from the members of the IETF 6lo working group. Great thanks to Samita Chakrabarti and Gabriel Montenegro for their feedback and support in connecting the IEEE and ITU-T sides. Authors thank Scott Mansfield, Ralph Droms, Pat Kinney for their guidance in the liaison process. Authors wish to thank Stefano Galli, Thierry Lys, Yizhou Li and Yuefeng Wu for their valuable comments and contributions.

9. References

9.1. Normative References

[I-D.ietf-6lo-rfc6775-update]

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

[I-D.ietf-roll-aodv-rpl]

Anamalamudi, S., Zhang, M., Perkins, C., Anand, S., and B. Liu, "Asymmetric AODV-P2P-RPL in Low-Power and Lossy Networks (LLNs)", draft-ietf-roll-aodv-rpl-05 (work in progress), October 2018.

- [IEEE_1901.1]
IEEE-SA Standards Board, "Standard for Medium Frequency (less than 15 MHz) Power Line Communications for Smart Grid Applications", IEEE 1901.1, May 2018, <<http://sites.ieee.org/sagroups-1901-1>>.
- [IEEE_1901.2]
IEEE-SA Standards Board, "IEEE Standard for Low-Frequency (less than 500 kHz) Narrowband Power Line Communications for Smart Grid Applications", IEEE 1901.2, October 2013, <<https://standards.ieee.org/findstds/standard/1901.2-2013.html>>.
- [ITU-T_G.9903]
International Telecommunication Union, "Narrowband orthogonal frequency division multiplexing power line communication transceivers for G3-PLC networks", ITU-T G.9903, February 2014, <<https://www.itu.int/rec/T-REC-G.9903>>.
- [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>>.
- [RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, DOI 10.17487/RFC2464, December 1998, <<https://www.rfc-editor.org/info/rfc2464>>.
- [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>>.
- [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>>.

- [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>>.
- [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>>.

9.2. Informative References

- [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.
- [IEEE_1901.2a] IEEE-SA Standards Board, "IEEE Standard for Low-Frequency (less than 500 kHz) Narrowband Power Line Communications for Smart Grid Applications - Amendment 1", IEEE 1901.2a, September 2015, <<https://standards.ieee.org/findstds/standard/1901.2a-2015.html>>.
- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <<https://www.rfc-editor.org/info/rfc3315>>.
- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, DOI 10.17487/RFC3972, March 2005, <<https://www.rfc-editor.org/info/rfc3972>>.
- [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>>.

- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, DOI 10.17487/RFC4941, September 2007, <<https://www.rfc-editor.org/info/rfc4941>>.
- [RFC5535] Bagnulo, M., "Hash-Based Addresses (HBA)", RFC 5535, DOI 10.17487/RFC5535, June 2009, <<https://www.rfc-editor.org/info/rfc5535>>.
- [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>>.
- [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>>.

Authors' Addresses

Jianqiang Hou
Huawei Technologies
101 Software Avenue,
Nanjing 210012
China

Email: [houjianqiang@huawei.com](mailto:hujianqiang@huawei.com)

Bing Liu
Huawei Technologies
No. 156 Beiqing Rd. Haidian District,
Beijing 100095
China

Email: remy.liubing@huawei.com

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

Email: yghong@etri.re.kr

Xiaojun Tang
State Grid Electric Power Research Institute
19 Chengxin Avenue
Nanjing 211106
China

Email: itc@sgepri.sgcc.com.cn

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

Email: charliep@computer.org

6Lo Working Group
Internet-Draft
Intended status: Informational
Expires: September 12, 2019

Y-G. Hong
ETRI
C. Gomez
UPC
Y-H. Choi
ETRI
AR. Sangi
Huaiyin Institute of Technology
T. Aanstoot
Modio AB
S. Chakrabarti
March 11, 2019

IPv6 over Constrained Node Networks (6lo) Applicability & Use cases
draft-ietf-6lo-use-cases-06

Abstract

This document describes the applicability of IPv6 over constrained node networks (6lo) and provides practical deployment examples. In addition to IEEE 802.15.4, various link layer technologies such as ITU-T G.9959 (Z-Wave), BLE, DECT-ULE, MS/TP, NFC, PLC (IEEE 1901.2), and IEEE 802.15.4e (6tisch) are used as examples. The document targets an audience who like to understand and evaluate running end-to-end IPv6 over the constrained node networks connecting devices to each other or to other devices on the Internet (e.g. cloud infrastructure).

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 September 12, 2019.

Copyright Notice

Copyright (c) 2019 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. Conventions and Terminology	4
3. 6lo Link layer technologies and possible candidates	4
3.1. ITU-T G.9959 (specified)	4
3.2. Bluetooth LE (specified)	4
3.3. DECT-ULE (specified)	5
3.4. MS/TP (specified)	5
3.5. NFC (specified)	6
3.6. PLC (specified)	7
3.7. IEEE 802.15.4e (specified)	7
3.8. Comparison between 6lo Link layer technologies	8
4. 6lo Deployment Scenarios	9
4.1. jupiternetwork in Smart Grid using 6lo in network layer	9
4.2. Wi-SUN usage of 6lo stacks	11
4.3. G3-PLC usage of 6lo in network layer	12
4.4. Netricity usage of 6lo in network layer	13
5. Design Space and Guidelines for 6lo Deployment	14
5.1. Design Space Dimensions for 6lo Deployment	14
5.2. Guidelines for adopting IPv6 stack (6lo/6LoWPAN)	16
6. 6lo Use Case Examples	17
7. IANA Considerations	18
8. Security Considerations	18
9. Acknowledgements	18
10. References	19
10.1. Normative References	19
10.2. Informative References	21
Appendix A. Other 6lo Use Case Examples	23
A.1. Use case of ITU-T G.9959: Smart Home	23
A.2. Use case of DECT-ULE: Smart Home	24
A.3. Use case of MS/TP: Building Automation Networks	25
A.4. Use case of NFC: Alternative Secure Transfer	25

A.5. Use case of PLC: Smart Grid	26
A.6. Use case of IEEE 802.15.4e: Industrial Automation	27
Authors' Addresses	27

1. Introduction

Running IPv6 on constrained node networks has different features from general node networks due to the characteristics of constrained node networks such as small packet size, short link-layer address, low bandwidth, network topology, low power, low cost, and large number of devices [RFC4919][RFC7228]. For example, some IEEE 802.15.4 link layers have a frame size of 127 octets and IPv6 requires the layer below to support an MTU of 1280 bytes, therefore an appropriate fragmentation and reassembly adaptation layer must be provided at the layer below IPv6. Also, the limited size of IEEE 802.15.4 frame and low energy consumption requirements make the need for header compression. The IETF 6LoPWAN (IPv6 over Low powerWPAN) working group published an adaptation layer for sending IPv6 packets over IEEE 802.15.4 [RFC4944], which includes a compression format for IPv6 datagrams over IEEE 802.15.4-based networks [RFC6282], and Neighbor Discovery Optimization for 6LoPWAN [RFC6775].

As IoT (Internet of Things) services become more popular, IPv6 over various link layer technologies 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), Power Line Communication (PLC), and IEEE 802.15.4e (TSCH), have been defined at [IETF_6lo] working group. IPv6 stacks for constrained node networks use a variation of the 6LoWPAN stack applied to each particular link layer technology.

In the 6LoPWAN working group, the [RFC6568], "Design and Application Spaces for 6LoWPANs" was published and it describes potential application scenarios and use cases for low-power wireless personal area networks. Hence, this 6lo applicability document aims to provide guidance to an audience who are new to IPv6-over-low-power networks concept and want to assess if variance of 6LoWPAN stack [6lo] can be applied to the constrained layer two (L2) network of their interest. This 6lo applicability document puts together various design space dimensions such as deployment, network size, power source, connectivity, multi-hop communication, traffic pattern, security level, mobility, and QoS requirements etc. In addition, it describes a few set of 6LoPWAN application scenarios and practical deployment as examples.

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

- o 6lo applicability and use cases MAY be uniquely different from those of 6LoWPAN defined for IEEE 802.15.4.
- o It SHOULD cover various IoT related wire/wireless link layer technologies providing practical information of such technologies.
- o A general guideline on how the 6LoWPAN stack can be modified for a given L2 technology.
- o Example use cases and practical deployment examples.

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 [RFC2119].

3. 6lo Link layer technologies and possible candidates

3.1. ITU-T G.9959 (specified)

The ITU-T G.9959 Recommendation [G.9959] targets low-power Personal Area Networks (PANs), 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.

3.2. Bluetooth LE (specified)

Bluetooth LE was introduced in Bluetooth 4.0, enhanced in Bluetooth 4.1, and developed even further in successive versions. Bluetooth SIG has also published 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 chipsets 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 such as phones, tablets and notebook computers. 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 [I-D.ietf-6lo-blemesh]

3.3. DECT-ULE (specified)

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 PHY layer operating at frequencies in the 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 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.

3.4. MS/TP (specified)

Master-Slave/Token-Passing (MS/TP) is a Medium Access Control (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 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]. MS/TP can be used for building automation networks.

3.5. NFC (specified)

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 [I-D.ietf-6lo-nfc]. NFC can be used for secure transfer in healthcare services.

3.6. PLC (specified)

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

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	<15MHz	PLC-IoT	10Mbps	2000m
IEEE1901.2	<500kHz	Narrowband	200Kbps	3000m

Table 1: Some Available Open Standards in PLC

[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 OFDM (Orthogonal Frequency Division Multiplexing) modulation.

[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 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.

3.7. IEEE 802.15.4e (specified)

The Time Slotted Channel Hopping (TSCH) mode was introduced in the IEEE 802.15.4-2015 standard. In a TSCH network, all nodes are synchronized. Time is sliced up into timeslots. The duration of a timeslot, typically 10ms, is large enough for a node to send a full-sized frame to its neighbor, and for that neighbor to send back an acknowledgment to indicate successful reception. Timeslots are grouped into one of more slotframes, which repeat over time.

All the communication in the network is orchestrated by a communication schedule which indicates to each node what to do in each of the timeslots of a slotframe: transmit, listen or sleep. The communication schedule can be built so that the right amount of link-layer resources (the cells in the schedule) are scheduled to satisfy the communication needs of the applications running on the network, while keeping the energy consumption of the nodes very low. Cells can be scheduled in a collision-free way, introducing a high level of determinism to the network.

A TSCH network exploits channel hopping: subsequent packet exchanges between neighbor nodes are done on a different frequency. This means that, if a frame isn't received, the transmitter node will re-transmit the frame on a different frequency. The resulting "channel hopping" efficiently combats external interference and multi-path fading.

The main benefits of IEEE 802.15.4 TSCH are:

- ultra high reliability. Off-the-shelf commercial products offer over 99.999% end-to-end reliability.
- ultra low-power consumption. Off-the-shelf commercial products offer over a decade of battery lifetime.
- 6TiSCH at IETF defines communications of TSCH network and it uses 6LoWPAN stack [RFC7554].

IEEE 802.15.4e can be used for industrial automation.

3.8. Comparison between 6lo Link layer technologies

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

	Z-Wave	BLE	DECT-ULE	MS/TP	NFC	PLC	TSCH
Usage	Home Auto-mation	Interact w/ Smart Phone	Meter Reading	Building Auto-mation	Health-care Service	Smart Grid	Industrial Aut-mation
Topology & Subnet	L2-mesh or L3-mesh	Star & Mesh	Star No mesh	MS/TP No mesh	P2P L2-mesh	Star Tree Mesh	Mesh
Mobility Reqmt	No	Low	No	No	Moderate	No	No
Security Reqmt	High + Privacy required	Parti-ally	High + Privacy required	High + Authen. required	High	High + Encrypt. required	High + Privacy required
Buffering Reqmt	Low	Low	Low	Low	Low	Low	Low
Latency, QoS Reqmt	High	Low	Low	High	High	Low	High
Data Rate	Infrequ-ent	Infrequ-ent	Infrequ-ent	Frequent	Small	Infrequ-ent	Infrequ-ent
RFC # or Draft	RFC7428	RFC7668	RFC8105	RFC8163	draft-ietf-6lo-nfc	draft-ietf-6lo-plc	RFC7554

Table 2: Comparison between 6lo Link layer technologies

4. 6lo Deployment Scenarios

4.1. jupitermesh in Smart Grid using 6lo in network layer

jupiterMesh is a multi-hop wireless mesh network specification designed mainly for deployment in large geographical areas. Each subnet in jupiterMesh is able to cover an entire neighborhood with thousands of nodes consisting of IPv6-enabled routers and end-points

(e.g. hosts). Automated network joining and load balancing allows a seamless deployment of a large number of subnets.

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

- o Automated meter reading
- o Distribution Automation (DA)
- o Demand-side management (DSM)
- o Demand-side response (DSR)
- o Power outage reporting
- o Street light monitoring and control
- o Transformer load management
- o EV charging coordination
- o Energy theft
- o Parking space locator

jupiterMesh specification is based on the following technologies:

- o The PHY layer is based on IEEE 802.15.4 SUN specification [IEEE 802.15.4-2015], supporting multiple operating modes for deployment in different regulatory domains and deployment scenarios in terms of density and bandwidth requirements. jupiterMesh supports bit rates from 50 kbps to 800 kbps, frame size up to 2048 bytes, up to 11 different RF bands and 3 modulation types (i.e., FSK, OQPSK and OFDM).
- o The MAC layer is based on IEEE 802.15.4 TSCH specification [IEEE 802.15.4-2015]. With frequency hopping capability, TSCH MAC supports scheduling of dedicated timeslot enabling bandwidth management and QoS.
- o The security layer consists of a certificate-based (i.e. X.509) network access authentication using EAP-TLS, with IEEE 802.15.9-based KMP (Key Management Protocol) transport, and PANA and link layer encryption using AES-128 CCM as specified in IEEE 802.15.4-2015 [IEEE 802.15.4-2015].

- o Address assignment and network configuration are specified using DHCPv6 [RFC3315]. Neighbor Discovery (ND) [RFC6775] and stateless address auto-configuration (SLAAC) are not supported.
- o The network layer consists of IPv6, ICMPv6 and 6lo/6LoPWAN header compression [RFC6282]. Multicast is supported using MPL. Two domains are supported, a delay sensitive MPL domain for low latency applications (e.g. DSM, DSR) and a delay insensitive one for less stringent applications (e.g. OTA file transfers).
- o The routing layer uses RPL [RFC6550] in non-storing mode with the MRHOF objective function based on the ETX metric.

4.2. Wi-SUN usage of 6lo stacks

Wireless Smart Ubiquitous Network (Wi-SUN) is a technology based on the IEEE 802.15.4g standard. Wi-SUN networks support star and mesh topologies, as well as hybrid star/mesh deployments, but 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.

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:

- o Advanced Metering Infrastructure (AMI)
- o Distribution Automation
- o Home Energy Management
- o Infrastructure Management
- o Intelligent Transportation Systems
- o Smart Street Lighting
- o Agriculture
- o Structural health (bridges, buildings etc)
- o Monitoring and Asset Management
- o Smart Thermostats, Air Conditioning and Heat Controls
- o 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. Examples include from meter to outdoor access point/router for AMI and DR, or between switches for DA. However, nothing in the profile restricts it to outdoor use. It has the following features;

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

In the Wi-SUN FAN specification, adaptation layer based on 6lo and IPv6 network layer are described. So, IPv6 protocol suite including TCP/UDP, 6lo Adaptation, Header Compression, DHCPv6 for IP address management, Routing using RPL, ICMPv6, and Unicast/Multicast forwarding is utilized.

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

G3-PLC [G3-PLC] is a narrow-band PLC technology that is based on ITU-T G.9903 Recommendation [G.9903]. G3-PLC supports multi-hop mesh network, 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 NB-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:

- o Smart Metering

- o Vehicle-to-Grid Communication
- o Demand Response (DR)
- o Distribution Automation
- o Home/Building Energy Management Systems
- o Smart Street Lighting
- o Advanced Metering Infrastructure (AMI) backbone network
- o Wind/Solar Farm Monitoring

In the G3-PLC specification, the 6lo adaptation layer utilizes the 6LoWPAN functions (e.g. header compression, fragmentation and reassembly) so as to enable IPv6 packet transmission. LOADng, which is a lightweight variant of AODV, is applied as the mesh-under routing protocol in G3-PLC networks. Address assignment and network configuration are based on the bootstrapping protocol specified in ITU-T G.9903. The network layer consists of IPv6 and ICMPv6 while the transport protocol UDP is used for data transmission.

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 1901.2 Low-Frequency Narrow-Band 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:

- o Utility grid modernization
- o Distribution automation
- o Meter-to-Grid connectivity
- o Micro-grids
- o Grid sensor communications
- o Load control

- o Demand response
- o Net metering
- o Street Lighting control
- o Photovoltaic panel monitoring

Netricity system architecture is based on the PHY and MAC layers of IEEE 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 layer 3 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. Design Space and Guidelines for 6lo Deployment

5.1. 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:

- o Deployment/Bootstrapping: 6lo nodes can be connected randomly, or in an organized manner. The bootstrapping has different characteristics for each link layer technology.
- o 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.
- o 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.
- o Multi-link subnet, single subnet: The selection of multi-link subnet and single subnet depends on connectivity and the number of 6lo nodes.

- o 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.
- o 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.
- o 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.
- o 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.
- o 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.
- o Reliability and QoS: Some 6lo use case requires high reliability, for example real-time service or health-related services.
- o 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.
- o Security Bootstrapping: Without the external operations, 6lo nodes must have the security bootstrapping mechanism.
- o 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.
- o 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.
- o 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.

5.2. Guidelines for adopting IPv6 stack (6lo/6LoWPAN)

The following guideline targets new candidate constrained L2 technologies that may be considered for running modified 6LoWPAN stack on top. The modification of 6LoWPAN stack should be based on the following:

- o **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. Whether the device is capable of forming IPv6 Link-local and global addresses, 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 address, 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.
- o **MTU Considerations:** The deployment SHOULD consider their need for maximum transmission unit (MTU) of a packet 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 6LoWPAN layer may skip supporting 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.
- o **Mesh or L3-Routing:** 6LoWPAN specifications do provide mechanisms to support for mesh routing at L2. [RFC6550] defines layer three (L3) routing for low power lossy networks using directed graphs. 6LoWPAN is routing protocol agnostic and other L2 or L3 routing protocols can be run using a 6LoWPAN stack.
- o **Address Assignment:** 6LoWPAN requires that IPv6 Neighbor Discovery for low power networks [RFC6775] be used for autoconfiguration of stateless IPv6 address assignment. Considering the energy sensitive networks [RFC6775] makes optimization from classical

IPv6 ND [RFC4861] protocol. It is the responsibility of the deployment to ensure unique global IPv6 addresses for the Internet connectivity. For local-only connectivity IPv6 ULA may be used. [RFC6775] specifies the 6LoWPAN border router(6LBR) which is responsible for prefix assignment to the 6lo/6LoWPAN network. 6LBR can be connected to the Internet or Enterprise network via its 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 autoconfiguration of IPv6 addressing due to regulatory and business reasons and may choose to offer a separate address assignment service.

- o Header Compression: IPv6 header compression [RFC6282] is a vital part of IPv6 over low power communication. Examples of header compression for different link-layers specifications are found in [RFC7668], [RFC8163], [RFC8105]. A generic header compression technique is specified in [RFC7400].
- o 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 [ace] and [core] should be consulted for application and transport level security. 6lo working group is working on address authentication [6lo-ap-nd] 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.
- o 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.

6. 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 clause, one 6lo use case example of Bluetooth LE (Smartphone-Based Interaction with Constrained Devices) is described. Other 6lo use case examples are described in Appendix.

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. 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.

7. IANA Considerations

There are no IANA considerations related to this document.

8. Security Considerations

Security considerations are not directly applicable to this document. The use cases will use the security requirements described in the protocol specifications.

9. Acknowledgements

Carles Gomez has been funded in part by the Spanish Government through the Jose Castillejo CAS15/00336 grant, and through the TEC2016-79988-P grant. 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, and Jianqiang HOU 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. 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.

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, <<https://www.rfc-editor.org/info/rfc2119>>.
- [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>>.
- [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>>.
- [RFC7554] Watteyne, T., Ed., Palattella, M., and L. Grieco, "Using IEEE 802.15.4e Time-Slotted Channel Hopping (TSCH) in the Internet of Things (IoT): Problem Statement", RFC 7554, DOI 10.17487/RFC7554, May 2015, <<https://www.rfc-editor.org/info/rfc7554>>.
- [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>>.
- [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>>.
- [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>>.

10.2. Informative References

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

- [I-D.ietf-6lo-nfc]
Choi, Y., Hong, Y., Youn, J., Kim, D., and J. Choi,
"Transmission of IPv6 Packets over Near Field
Communication", draft-ietf-6lo-nfc-13 (work in progress),
February 2019.
- [I-D.ietf-roll-aodv-rpl]
Anamalamudi, S., Zhang, M., Perkins, C., Anand, S., and B.
Liu, "Asymmetric AODV-P2P-RPL in Low-Power and Lossy
Networks (LLNs)", draft-ietf-roll-aodv-rpl-06 (work in
progress), March 2019.
- [I-D.ietf-6tisch-6top-sfx]
Dujovne, D., Grieco, L., Palattella, M., and N. Accettura,
"6TiSCH Experimental Scheduling Function (SFX)", draft-
ietf-6tisch-6top-sfx-01 (work in progress), March 2018.
- [I-D.ietf-6lo-blemesh]
Gomez, C., Darroudi, S., Savolainen, T., and M. Spoerk,
"IPv6 Mesh over BLUETOOTH(R) Low Energy using IPSP",
draft-ietf-6lo-blemesh-04 (work in progress), January
2019.
- [I-D.satish-6tisch-6top-sf1]
Anamalamudi, S., Liu, B., Zhang, M., Sangi, A., Perkins,
C., and S. Anand, "Scheduling Function One (SF1): hop-by-
hop Scheduling with RSVP-TE in 6tisch Networks", draft-
satish-6tisch-6top-sf1-04 (work in progress), October
2017.
- [I-D.ietf-6lo-plc]
Hou, J., Liu, B., Hong, Y., Tang, X., and C. Perkins,
"Transmission of IPv6 Packets over PLC Networks", draft-
ietf-6lo-plc-00 (work in progress), February 2019.
- [IETF_6lo]
"IETF IPv6 over Networks of Resource-constrained Nodes
(6lo) working group",
<<https://datatracker.ietf.org/wg/6lo/charter/>>.
- [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](https://global.ihs.com/doc_detail.cfm?item_s_key=00032964)>.
- [G3-PLC] "G3-PLC Alliance", <<http://www.g3-plc.com/home/>>.

- [NETRICITY] "Netricity program in HomePlug Powerline Alliance", <<http://groups.homeplug.org/tech/Netricity>>.
- [G.9959] "International Telecommunication Union, "Short range narrow-band digital radiocommunication transceivers - PHY and MAC layer specifications", ITU-T Recommendation", January 2015.
- [G.9903] "International Telecommunication Union, "Narrowband orthogonal frequency division multiplexing power line communication transceivers for G3-PLC networks", ITU-T Recommendation", August 2017.
- [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 (work-in-progress), IEEE-SA Standards Board", <<http://sites.ieee.org/sagroups-1901-1/>>.
- [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>>.
- [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>.

Appendix A. Other 6lo Use Case Examples

A.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. a light is turned off). 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].

A.2. 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.

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.

A.3. 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.

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-site restrictions or hop-by-hop latency of many low cost wireless solutions.

A.4. Use case of NFC: Alternative Secure Transfer

According to applications, various secured data can be handled and transferred. Depending on security level of the data, methods for transfer 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. An LoWPAN Border Router (LBR) 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. At this moment, hidden hackers can overhear the data based on the LTE connection, but they cannot gather the personal data over the NFC connection.

A.5. Use case of PLC: Smart Grid

Smart grid concept is based on numerous operational and energy measuring sub-systems of an electric grid. It comprises of multiple administrative levels/segments to provide connectivity among these numerous components. Last mile connectivity is established over LV segment, whereas connectivity over electricity distribution takes place in HV segment.

Although other wired and wireless technologies are also used in Smart Grid (Advance Metering Infrastructure - AMI, Demand Response - DR, Home Energy Management System - HEMS, Wide Area Situational Awareness - WASA etc), PLC enjoys the advantage of existing (power conductor) medium and better reliable data communication. PLC is a promising wired communication technology in that the electrical power lines are already there and the deployment cost can be comparable to wireless technologies. The 6lo related scenarios lie in the low voltage 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 (IEEE1901.1) for WASA in Smart Grid

Many sub-systems of Smart Grid require low data rate and narrowband variant (IEEE1901.2) of PLC fulfils 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.

A.6. Use case of IEEE 802.15.4e: Industrial Automation

Typical scenario of Industrial Automation where sensor and actuators are connected through the time-slotted radio access (IEEE 802.15.4e). For that, there will be a point-to-point control signal exchange in between sensors and actuators to trigger the critical control information. In such scenarios, point-to-point traffic flows are significant to exchange the controlled information in between sensors and actuators within the constrained networks.

Example: Use of IEEE 802.15.4e for P2P communication in closed-loop application

AODV-RPL [I-D.ietf-roll-aodv-rpl] is proposed as a standard P2P routing protocol to provide the hop-by-hop data transmission in closed-loop constrained networks. Scheduling Functions i.e. SF0 [I-D.ietf-6tisch-6top-sfx] and SF1 [I-D.satish-6tisch-6top-sf1] is proposed to provide distributed neighbor-to-neighbor and end-to-end resource reservations, respectively for traffic flows in deterministic networks (6TiSCH).

The potential scenarios that can make use of the end-to-end resource reservations can be in health-care and industrial applications. AODV-RPL and SF0/SF1 are the significant routing and resource reservation protocols for closed-loop applications in constrained networks.

Authors' Addresses

Yong-Geun Hong
ETRI
161 Gajeong-Dong Yuseung-Gu
Daejeon 305-700
Korea

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

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

Email: carlesgo@entel.upc.edu

Younghwan Choi
ETRI
218 Gajeongno, Yuseong
Daejeon 305-700
Korea

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

Take Aanstoot
Modio AB
S:t Larsgatan 15, 582 24
Linkoping
Sweden

Email: take@modio.se

Samita Chakrabarti
San Jose, CA
USA

Email: samitac.ietf@gmail.com

6lo
Internet-Draft
Updates: 8505 (if approved)
Intended status: Standards Track
Expires: July 29, 2019

P. Thubert, Ed.
E. Levy-Abegnoli
Cisco Systems
January 25, 2019

IPv6 Neighbor Discovery Unicast Lookup
draft-thubert-6lo-unicast-lookup-00

Abstract

This document updates RFC 8505 in order to enable unicast address lookup from a 6LoWPAN Border Router acting as an Address Registrar.

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 July 29, 2019.

Copyright Notice

Copyright (c) 2019 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	3
2.1. BCP 14	3
2.2. References	4
2.3. New Terms	4
2.4. Acronym Definitions	4
3. Overview	5
4. Updating RFC 8505	7
4.1. Extended Neighbor Discovery Options and Messages	7
4.1.1. Extending the Capability Indication Option	8
4.1.2. New Code Prefix for Address Mapping Messages	8
4.1.3. New ARO Status	8
4.2. Address Mapping Messages	9
4.3. IPv6 ND-based Address Lookup	10
5. Backward Compatibility	10
6. Security Considerations	10
7. IANA Considerations	10
7.1. ICMP Codes	11
7.2. New ARO Status values	11
7.3. New 6LoWPAN Capability Bits	12
8. Acknowledgments	12
9. References	12
9.1. Normative References	12
9.2. Informative References	13
Authors' Addresses	14

1. Introduction

[RFC8505] defines the Routing Registrar and extends [RFC6775] to use a 6LoWPAN Border Router (6LBR) as a central service for Address Registration and duplicate detection amongst Routing Registrars and possibly individual Nodes that access it directly.

[I-D.ietf-6lo-backbone-router] introduces the Backbone Router (6BBR) as a Routing Registrar that performs IPv6 ND [RFC4861] [RFC4862] proxy operation between IPv6 Nodes on a federating Backbone Link and Registering Nodes attached to a LowPower Lossy Networks (LLNs) that register their addresses to the 6BBR. The federated links form a Multilink Subnet (MLSN).

The 6BBRs may exchange Extended Duplicate Address Messages (EDAR and EDAC) [RFC8505] to register the proxied addresses on behalf of the Registering Nodes to the 6LBR. The Registration Ownership Verifier (ROVR) field in the EDAR and EDAC messages is used to correlate attempts to register the same address and to detect duplications. The ROVR can also be used as a proof-of-ownership (see

[I-D.ietf-6lo-ap-nd]) to protect the Registered address against theft and impersonation attacks (more in [I-D.bi-savi-wlan]). Conflicting registrations to different 6BBRs for the same Registered address are resolved using the TID field, which creates a temporal order and enables to recognize the freshest registration.

With [I-D.ietf-6lo-backbone-router], the Link Layer address (LLA) that the 6BBR advertises for a Registered address on behalf of the Registered Node over the Backbone can belong to the Registering Node; in that case, the 6BBR acts as a Bridging Proxy and bridges the unicast packets. Alternatively, the LLA can be that of the 6BBR on the Backbone interface, in which case the 6BBR acts as a Routing Proxy, that receives the unicast packets at Layer-3 and routes them. The 6BBR signals that LLA in a Source LLA Option (SLLAO) in the EDAR messages to the 6LBR, and the 6LBR responds with a Target LLA Option (TLLAO) that indicates the LLA associated to the current registration.

It results that the 6LBR is capable of providing the LLA mapping for any address that was proactively registered with an SLLAO. This draft defines the protocol elements and the operations to try a unicast lookup with the 6LBR. This may save a reactive IPv6 ND Neighbor Solicitation (NS) message, which is based on multicast and may be problematic in extensive wireless domains (see [I-D.ietf-mboned-ieee802-mcast-problems]) as well as in large switched fabrics.

The registration and lookup services that the 6LBR provides do not have to be limited to 6BBRs and are available to any node that supports [RFC8505] and [I-D.ietf-6lo-backbone-router] to register an address, and / or this specification to resolve a mapping. The services are available on-link using an IPv6 NDP NS and off-link using a new variation of the Extended Duplicate Address messages called Address Mapping Messages. The policy and security settings that allow the access to the 6LBR are out of scope.

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. References

This document uses terms and concepts that are discussed in:

- o "Neighbor Discovery for IP version 6" [RFC4861] and "IPv6 Stateless address Autoconfiguration" [RFC4862],
- o Neighbor Discovery Optimization for Low-Power and Lossy Networks [RFC6775], as well as
- o "Registration Extensions for 6LoWPAN Neighbor Discovery" [RFC8505] and "IPv6 Backbone Router" [I-D.ietf-6lo-backbone-router].

2.3. New Terms

This document introduces the following terminology:

Address Mapping Request

An ICMP message with an ICMP type of 157 (DAR) and a Code Prefix of 1.

Address Mapping Confirm

An ICMP message with an ICMP type of 158 (DAC) and a Code Prefix of 1.

Address Registrar

The Address Registrar is an abstract database that is maintained by the 6LBR to store the state associated with its registrations.

Address Registration

An Address Registration is an abstract state associated to one registration, in other words one entry in the Address Registrar.

2.4. Acronym Definitions

This document uses the following acronyms:

6BBR: 6LoWPAN Backbone Router

6LBR: 6LoWPAN Border Router

6LR: 6LoWPAN Router

6CIO: Capability Indication Option

AMC: Address Mapping Confirmation

AMR: Address Mapping Request

ARO: Address Registration Option

DAC: Duplicate Address Confirmation

DAD: Duplicate Address Detection

DAR: Duplicate Address Request

EDAC: Extended Duplicate Address Confirmation

EDAR: Extended Duplicate Address Request

DODAG: Destination-Oriented Directed Acyclic Graph

LLN: Low-Power and Lossy Network

NA: Neighbor Advertisement

NCE: Neighbor Cache Entry

ND: Neighbor Discovery

NS: Neighbor Solicitation

ROVR: Registration Ownership Verifier

RA: Router Advertisement

RS: Router Solicitation

TID: Transaction ID

3. Overview

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.

A collection of IPv6 Nodes are present on the Backbone and use IPv6 ND [RFC4861][RFC4862] procedures for DAD and Lookup.

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].

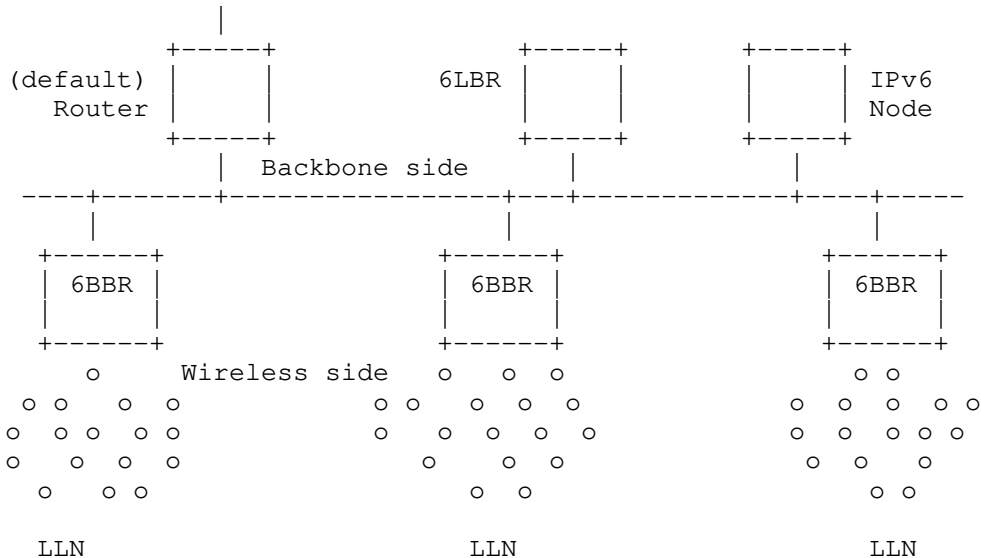


Figure 1: Backbone Link and 6LBR

A 6LBR provides registration services for the purpose of proactive IPv6 ND and maintains a registry of the active registrations as an abstract data structure called an Address Registrar. An entry in the Address Registrar is called an "Address Registration".

The Address Registration retains:

- o the value for the ROVR associated to the registration, the current value of the TID, and the remaining Lifetime.
- o a list of LLAs that are associated with the IPv6 address and can be used in a TLLAO as a response to a lookup.

Examples where more than one address may be available include the case of an anycast address and the case of an LLN address that is proxied by more than one 6BBR.

Unless otherwise configured, a 6LBR does the following:

- o The 6LBR maintains an entry in the Address Registrar for any type of unicast and anycast addresses including those with link-local scope.

- o Based on that entry, it provides duplicate avoidance services within the scope of its Address Registrar.
- o The 6LBR also provides address lookup services for the Registered Address using unicast ICMPv6 DAR and DAC-based Address Mapping messages.

The Address Mapping messages can be exchanged using global unicast addresses as source and destination addresses, so they can be used for both on-link and off-link queries. NS and NA messages may also be used, but in that case the unicast source and destination addresses are link-local addresses and the 6LBR must be on-link.

The 6LBR proactive operations may coexist on the Backbone with reactive IPv6 ND [RFC4861][RFC4862] that rely on multicast for Duplicate Address Detection (DAD) and Address Lookup. Nodes that support this specification operate with the 6LBR before attempting the reactive operation, which may be avoided if the 6LBR is conclusive, either detecting a duplication or returning a mapping.

4. Updating RFC 8505

This specification leverages the capability to insert IPv6 ND options in the EDAR and EDAC messages that was introduced in [I-D.ietf-6lo-backbone-router].

It extends DAR and DAR ICMP messages for address lookup in Section 4.1.2 that use the same ICMP types as EDAR and EDAC but a different Code Prefix.

It also adds a new Status "Not Found" in Section 4.1.3) that indicates that the address being searched is not present in the Address Registrar.

A 6LBR signals itself by setting the "B" bit in the 6CIO of the RA messages that it generates [RFC8505]. This specification adds a new "A" bit in the 6CIO to indicate support of address mapping (see Section 4.1.1).

4.1. Extended Neighbor Discovery Options and Messages

This specification does not introduce new options; it modifies existing options and updates the associated behaviors.

4.1.1.1. Extending the Capability Indication Option

This specification defines a new capability bit for use in the 6CIO, as defined by [RFC7400] and extended in[RFC8505] for use in IPv6 ND messages.

The new "A" bit indicates that the 6LBR provides address mapping services per this specification.




Figure 2: New Capability Bits in the 6CIO

Option Fields:

Type: 36

A: The 6LBR provides address mapping services.

4.1.1.2. New Code Prefix for Address Mapping Messages

The Extended Duplicate Address messages share a common base format defined in section 4.2 of [RFC8505], with the ICMP type respectively set to 157 and 158 that is inherited from the DAR and DAC messages defined in section 4.4 of [RFC6775]. The ICMP Code is split in two 4-bit fields, the Code Prefix and the Code Suffix, and the only Code Prefix defined in [RFC8505] is 0, signaling a DAD.

The Address Mapping messages use the same values for the ICMP Type as the corresponding Extended Duplicate Address messages. This specification adds the Code Prefix of 1 to signal Address Mapping. ICMP messages with the ICMP type set to 157 or 158, and a Code Prefix of 1 are thus respectively an Address Mapping Request (AMR) and an Address Mapping Confirm (AMC).

4.1.1.3. New ARO Status

The Extended Address Registration Option (EARO) is defined in section 4.1 of [RFC8505]. It contains a Status field that is common with the EDAR and EDAC messages defined in section 4.2 of [RFC8505].

This specification defines a new Status "Not Found" as indicated in Table 1

Value	Description
0..10	As defined in [RFC6775] and [RFC8505].
11	Not Found: The address is not present in the Address Registrar (value to be confirmed by IANA)

Table 1: EARO Status

The Status of "Not Found" can be used in an NA(EARO) and in an AMC messages as a response to an address lookup operation.

4.2. Address Mapping Messages

A 6LBR signals that support by setting the "B" bit in the 6CIO of the RA messages that it generates. A 6LBR that supports this specification MUST also set the "A" bit, indicating support of the Address Mapping messages for address lookup.

In the Address Mapping flow, the querier IPv6 Node uses an AMR message, which is characterized by an ICMPv6 Type of 157 and a Code Prefix of 1. When used on-link, the AMR message SHOULD carry a SLLAO indicating the LLA of the querier. The Code Suffix MUST be set to 0 indicating a ROVR Length of 64 bits. The ROVR, TID and Lifetime fields MUST be set to 0 and ignored by the receiver.

The 6LBR MUST respond with an AMC message, which is characterized by an ICMPv6 Type of 158 and a Code Prefix of 1.

- o If the address is not present in the Address Registrar then the 6LBR MUST set the status to "Not Found". The Code Suffix MUST be set to 0 indicating a ROVR Length of 64 bits. The ROVR, TID and Lifetime fields MUST be set to 0 and ignored by the receiver.
- o Else if the address is present in the Address Registrar then the AMC fields MUST be set from the ROVR, TID and remaining Lifetime values in the Address Registration and the Status MUST be set to 0.
- o If at least one LLA is found in the Address Registration, then the 6LBR MUST place one in a TLLAO option in the AMC message.

The AMC is sent unicast the 6LBR to the querier.

4.3. IPv6 ND-based Address Lookup

A 6LBR that is deployed on-link SHOULD provide NS/NA-based services. It signals that support by setting the "L" bit in the 6CIO of the RA messages that it generates, indicating that it is a 6LR [RFC8505].

A 6LBR thus typically sets the "A", the "B", and the "L" bits when attached to a Backbone Link that it serves, as illustrated in Figure 1. In that case, the IPv6 Nodes and 6BBRs can use an NS/NA exchange with the 6LBR for both duplicate detection and lookup services.

The NS(Lookup) is sent unicast from link-local address of the querier to the link-local address of the 6LBR. It carries a SLLAO [RFC4861] and it MUST NOT carry an EARO option to avoid the confusion with a registration.

The 6LBR MUST respond with an NA message that contains an EARO.

- o If the address is not present in the Address Registrar then the 6LBR MUST set the status to "Not Found". The ROVR, TID and Lifetime fields MUST be set to 0 and ignored by the receiver.
- o Else if the address is present in the Address Registrar then the EARO fields MUST be set from the ROVR, TID and remaining Lifetime values in the Address Registration and the Status MUST be set to 0.
- o If at least one LLA is found in the Address Registration, then the 6LBR MUST place one in a TLLAO option in the NA message.

The NA is sent unicast from link-local address of the 6LBR to the link-local address of the querier.

5. Backward Compatibility

6. Security Considerations

This specification extends [RFC8505], and the security section of that document also applies to this document. In particular, the link layer SHOULD be sufficiently protected to prevent rogue access.

7. IANA Considerations

Note to RFC Editor, to be removed: please replace "This RFC" throughout this document by the RFC number for this specification once it is allocated.

IANA is requested to make a number of changes under the "Internet Control Message Protocol version 6 (ICMPv6) Parameters" registry, as follows.

7.1. ICMP Codes

IANA is requested to create 2 new subregistries of the ICMPv6 "Code" Fields registry, which itself is a subregistry of the Internet Control Message Protocol version 6 (ICMPv6) Parameters for the ICMP codes.

The new subregistries relate to the ICMP type 157, Duplicate Address Request (shown in Table 2), and 158, Duplicate Address Confirmation (shown in Table 3), respectively. For those two ICMP types, the ICMP Code field is split into 2 subfields, the "Code Prefix" and the "Code". The new subregistries relate to the "Code Prefix" portion of the ICMP Code. The range of "Code Prefix" is 0..15 in all cases. The policy is "IETF Review" or "IESG Approval" [RFC8126] for both subregistries.

The new subregistries are to be initialized as follows:

Code Prefix	Meaning	Reference
0	Duplicate Address Detection	RFC 6775
1	Address Mapping	This RFC
2...15	Unassigned	

Table 2: New Code Prefixes for ICMP type 157 DAR message

Code Prefix	Meaning	Reference
0	Duplicate Address Detection	RFC 6775
1	Address Mapping	This RFC
2...15	Unassigned	

Table 3: New Code Prefixes for ICMP type 158 DAC message

7.2. New ARO Status values

IANA is requested to make additions to the Address Registration Option Status Values Registry as follows:

ARO Status	Description	Document
11	Not Found	This RFC

Table 4: New ARO Status values

7.3. New 6LoWPAN Capability Bits

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

Capability Bit	Description	Document
9	AM Support (A bit)	This RFC

Table 5: New 6LoWPAN Capability Bits

8. Acknowledgments

9. References

9.1. 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, <<https://www.rfc-editor.org/info/rfc2119>>.
- [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>>.
- [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>>.
- [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>>.
- [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>>.

9.2. Informative References

- [I-D.bi-savi-wlan]
Bi, J., Wu, J., Wang, Y., and T. Lin, "A SAVI Solution for WLAN", draft-bi-savi-wlan-16 (work in progress), November 2018.
- [I-D.ietf-6lo-ap-nd]
Thubert, P., Sethi, M., Struik, R., and B. Sarikaya, "Address Protected Neighbor Discovery for Low-power and Lossy Networks", draft-ietf-6lo-ap-nd-09 (work in progress), December 2018.
- [I-D.ietf-6lo-backbone-router]
Thubert, P., Perkins, C., and E. Levy-Abegnoli, "IPv6 Backbone Router", draft-ietf-6lo-backbone-router-10 (work in progress), January 2019.
- [I-D.ietf-mboned-ieee802-mcast-problems]
Perkins, C., McBride, M., Stanley, D., Kumari, W., and J. Zuniga, "Multicast Considerations over IEEE 802 Wireless Media", draft-ietf-mboned-ieee802-mcast-problems-04 (work in progress), November 2018.

[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)".

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

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

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