

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

P. Thubert, Ed.
Cisco
B. Sarikaya

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

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

Abstract

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

Status of This Memo

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

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

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

This Internet-Draft will expire on 1 November 2020.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

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

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

1. Introduction

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

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

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

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

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

2. Terminology

2.1. BCP 14

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

2.2. Additional References

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

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

2.3. Abbreviations

This document uses the following abbreviations:

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

3. Updating RFC 8505

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

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

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

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

4. New Fields and Options

4.1. New Crypto-ID

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

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

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

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

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

4.2. Updated EARO

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

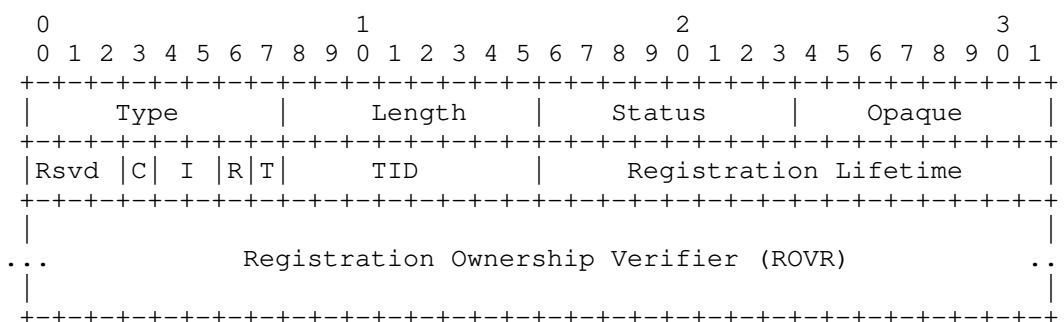


Figure 1: Enhanced Address Registration Option

Type: 33

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

Status: Defined in [RFC8505].

Opaque: Defined in [RFC8505].

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

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

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

TID: Defined in [RFC8505].

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

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

this specification does not define any new Status value.

4.3. Crypto-ID Parameters Option

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

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

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

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




Figure 2: Crypto-ID Parameters Option

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

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

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

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

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

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

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

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

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

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

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

4.4. NDP Signature Option

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

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

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

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

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




Figure 3: NDP signature Option

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

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

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

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

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

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

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

4.5. Extensions to the Capability Indication Option

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




Figure 4: New Capability Bit in the 6CIO

New Option Field:

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

The "A" flag is set by the 6LBR that serves the network and propagated by the 6LRs. It is typically turned on when all 6LRs are migrated to this specification.

5. Protocol Scope

The scope of the protocol specified here is a 6LoWPAN 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 5: 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 all the LLN links between the 6LR and the 6LBR are protected 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 ROVR associated to the address being registered upon the first registration and rejecting a registration with a different ROVR value. A 6LN can claim any address as long as it is the first to make that claim. After a successful registration, the 6LN becomes the owner of the registered address and the address is bound to the ROVR value in the 6LR/6LBR registry.

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

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

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

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

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

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

6.1. First Exchange with a 6LR

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

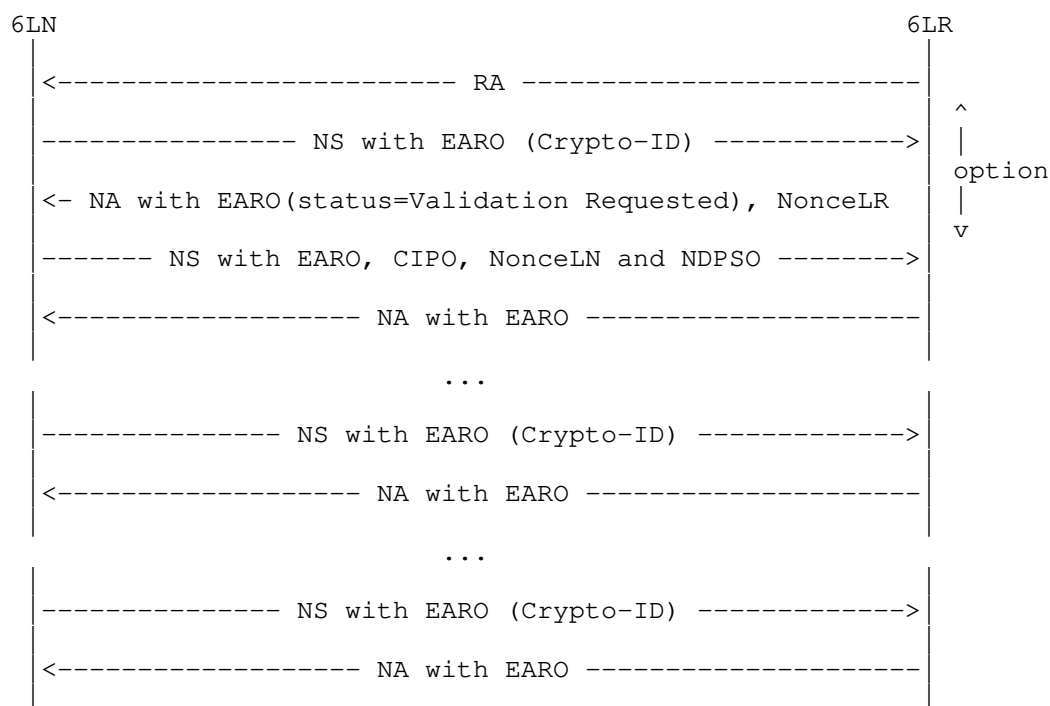


Figure 6: On-link Protocol Operation

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

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

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

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

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

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

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

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

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

6.2. NDPSO generation and verification

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

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

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

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

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

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

6.3. Multihop Operation

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

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

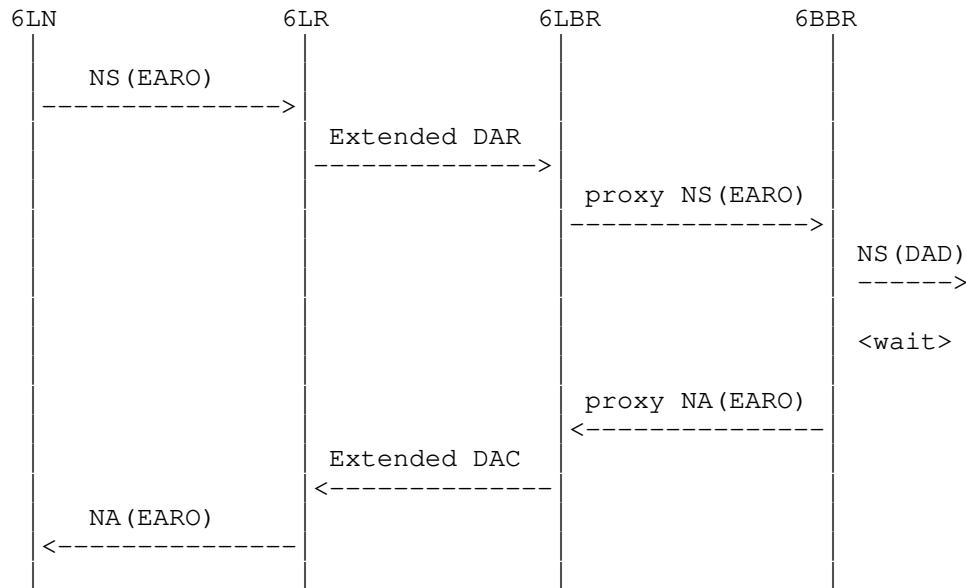


Figure 7: (Re-)Registration Flow

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

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

7. Security Considerations

7.1. Brown Field

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

7.2. Inheriting from RFC 3971

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

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

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

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

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

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

7.3. Related to 6LoWPAN ND

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

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

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

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

7.4. Compromised 6LR

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

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

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

7.5. ROVR Collisions

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

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

7.6. Implementation Attacks

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

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

7.7. Cross-Algorithm and Cross-Protocol Attacks

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

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

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

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

7.8. Public Key Validation

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

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

7.9. Correlating Registrations

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

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

8. IANA considerations

8.1. CGA Message Type

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

8.2. Crypto-Type Subregistry

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

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

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

Table 1: Crypto-Types

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

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

8.3. IPv6 ND option types

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

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

Table 2: New ND options

8.4. New 6LoWPAN Capability Bit

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

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

Table 3: New 6LoWPAN Capability Bit

9. Acknowledgments

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

10. Normative References

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

- [RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, DOI 10.17487/RFC3971, March 2005, <<https://www.rfc-editor.org/info/rfc3971>>.
- [RFC6234] Eastlake 3rd, D. and T. Hansen, "US Secure Hash Algorithms (SHA and SHA-based HMAC and HKDF)", RFC 6234, DOI 10.17487/RFC6234, May 2011, <<https://www.rfc-editor.org/info/rfc6234>>.
- [RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <<https://www.rfc-editor.org/info/rfc6775>>.
- [RFC7400] Bormann, C., "6LoWPAN-GHC: Generic Header Compression for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 7400, DOI 10.17487/RFC7400, November 2014, <<https://www.rfc-editor.org/info/rfc7400>>.
- [RFC7748] Langley, A., Hamburg, M., and S. Turner, "Elliptic Curves for Security", RFC 7748, DOI 10.17487/RFC7748, January 2016, <<https://www.rfc-editor.org/info/rfc7748>>.
- [RFC8032] Josefsson, S. and I. Liusvaara, "Edwards-Curve Digital Signature Algorithm (EdDSA)", RFC 8032, DOI 10.17487/RFC8032, January 2017, <<https://www.rfc-editor.org/info/rfc8032>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8505] Thubert, P., Ed., Nordmark, E., Chakrabarti, S., and C. Perkins, "Registration Extensions for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Neighbor Discovery", RFC 8505, DOI 10.17487/RFC8505, November 2018, <<https://www.rfc-editor.org/info/rfc8505>>.
- [FIPS186-4] FIPS 186-4, "Digital Signature Standard (DSS), Federal Information Processing Standards Publication 186-4", US Department of Commerce/National Institute of Standards and Technology , July 2013.
- [SEC1] SEC1, "SEC 1: Elliptic Curve Cryptography, Version 2.0", Standards for Efficient Cryptography , June 2009.

11. Informative references

- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, DOI 10.17487/RFC3972, March 2005, <<https://www.rfc-editor.org/info/rfc3972>>.
- [BCP 106] Eastlake 3rd, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", BCP 106, RFC 4086, DOI 10.17487/RFC4086, June 2005, <<https://www.rfc-editor.org/info/rfc4086>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <<https://www.rfc-editor.org/info/rfc4862>>.
- [RFC4919] Kushalnagar, N., Montenegro, G., and C. Schumacher, "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals", RFC 4919, DOI 10.17487/RFC4919, August 2007, <<https://www.rfc-editor.org/info/rfc4919>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<https://www.rfc-editor.org/info/rfc4944>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<https://www.rfc-editor.org/info/rfc6282>>.
- [RFC7039] Wu, J., Bi, J., Bagnulo, M., Baker, F., and C. Vogt, Ed., "Source Address Validation Improvement (SAVI) Framework", RFC 7039, DOI 10.17487/RFC7039, October 2013, <<https://www.rfc-editor.org/info/rfc7039>>.
- [RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <<https://www.rfc-editor.org/info/rfc7217>>.

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

Appendix A. Requirements Addressed in this Document

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

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

Appendix B. Representation Conventions

B.1. Signature Schemes

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

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

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

B.2. Representation of ECDSA Signatures

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

B.3. Representation of Public Keys Used with ECDSA

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

B.4. Alternative Representations of Curve25519

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

General parameters (for all curve models):

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

Montgomery curve-specific parameters (for Curve25519):

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

Gv

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

Twisted Edwards curve-specific parameters (for Edwards25519):

a -1 (-0x01)

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

Gx

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

Gy 4/5

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

Weierstrass curve-specific parameters (for Wei25519):

a

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

b

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

GX

```
192986815395526992372618308347813179755449974442734273399095973346
52188435546
(=0x2aaaaaaaa aaaaaaaaa aaaaaaaaa aaaaaaaaa aaaaaaaaa aaaaaaaaa aaaaaaa
aad245a)
```

GY

```
147816194475895447910205935684099868872646061346164752889648818377
55586237401
```

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

Authors' Addresses

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

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

Behcet Sarikaya

Email: sarikaya@ieee.org

Mohit Sethi
Ericsson
FI-02420 Jorvas
Finland

Email: mohit@piuha.net

Rene Struik
Struik Security Consultancy

Email: rstruik.ext@gmail.com

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

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

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

Abstract

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

Status of This Memo

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

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

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

This Internet-Draft will expire on 24 September 2020.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text

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

Table of Contents

1. Introduction	2
2. Terminology	5
2.1. BCP 14	5
2.2. New Terms	5
2.3. Abbreviations	6
2.4. References	7
3. Overview	7
3.1. Updating RFC 6775 and RFC 8505	10
3.2. Access Link	11
3.3. Route-Over Mesh	13
3.4. The Binding Table	14
3.5. Primary and Secondary 6BBRs	15
3.6. Using Optimistic DAD	16
4. Multi-Link Subnet Considerations	17
5. Optional 6LBR serving the Multi-Link Subnet	17
6. Using IPv6 ND Over the Backbone Link	18
7. Routing Proxy Operations	20
8. Bridging Proxy Operations	21
9. Creating and Maintaining a Binding	22
9.1. Operations on a Binding in Tentative State	23
9.2. Operations on a Binding in Reachable State	24
9.3. Operations on a Binding in Stale State	25
10. Registering Node Considerations	26
11. Security Considerations	27
12. Protocol Constants	30
13. IANA Considerations	30
14. Acknowledgments	30
15. Normative References	30
16. Informative References	32
Appendix A. Possible Future Extensions	34
Appendix B. Applicability and Requirements Served	35
Authors' Addresses	37

1. Introduction

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

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

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

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

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

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

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

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

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

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

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

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

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

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

2. Terminology

2.1. BCP 14

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

2.2. New Terms

This document introduces the following terminology:

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

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

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

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

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

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

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

2.3. Abbreviations

This document uses the following abbreviations:

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

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

2.4. References

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

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

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

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

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

3. Overview

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

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

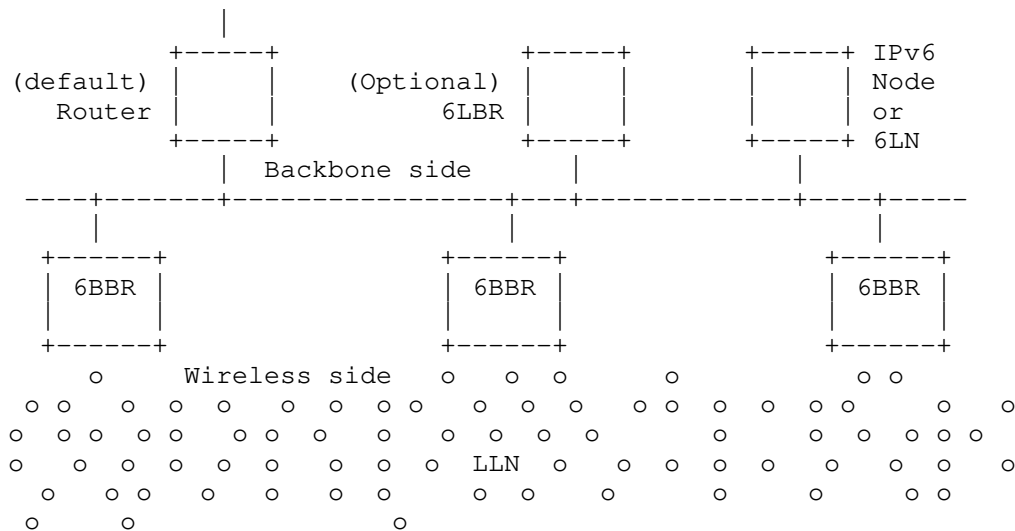


Figure 1: Backbone Link and Backbone Routers

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

The main features of a 6BBR are as follows:

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

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

Unless otherwise configured, a 6BBR does the following:

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

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

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

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

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

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

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

3.1. Updating RFC 6775 and RFC 8505

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

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

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

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

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

3.2. Access Link

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

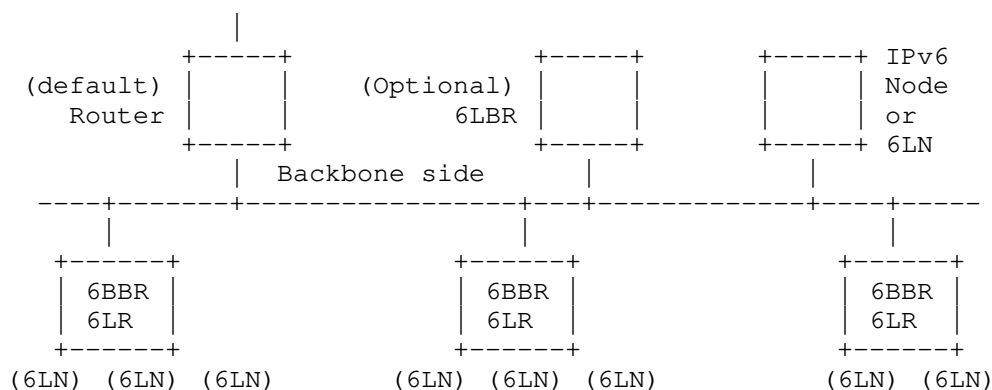


Figure 2: Access Link Use case

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

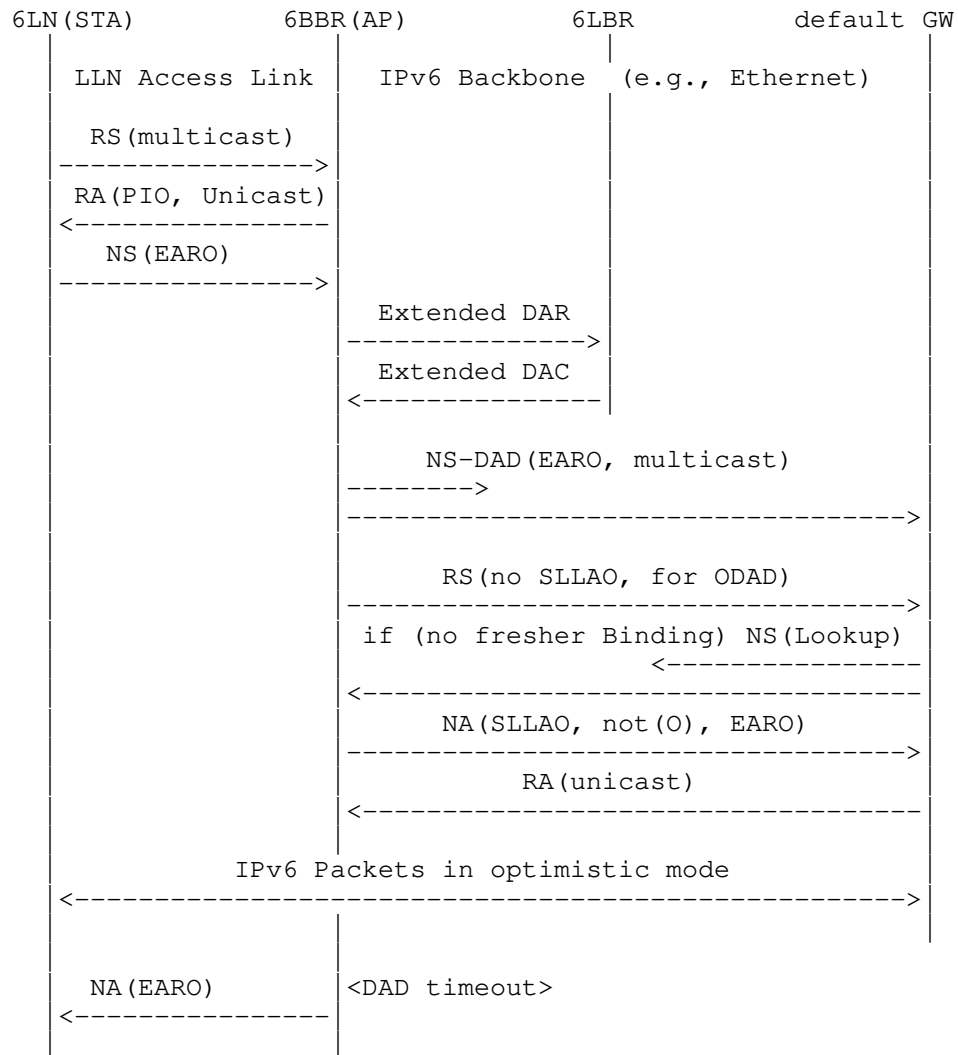


Figure 3: Initial Registration Flow to a 6BBR acting as Routing Proxy

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

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

3.3. Route-Over Mesh

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

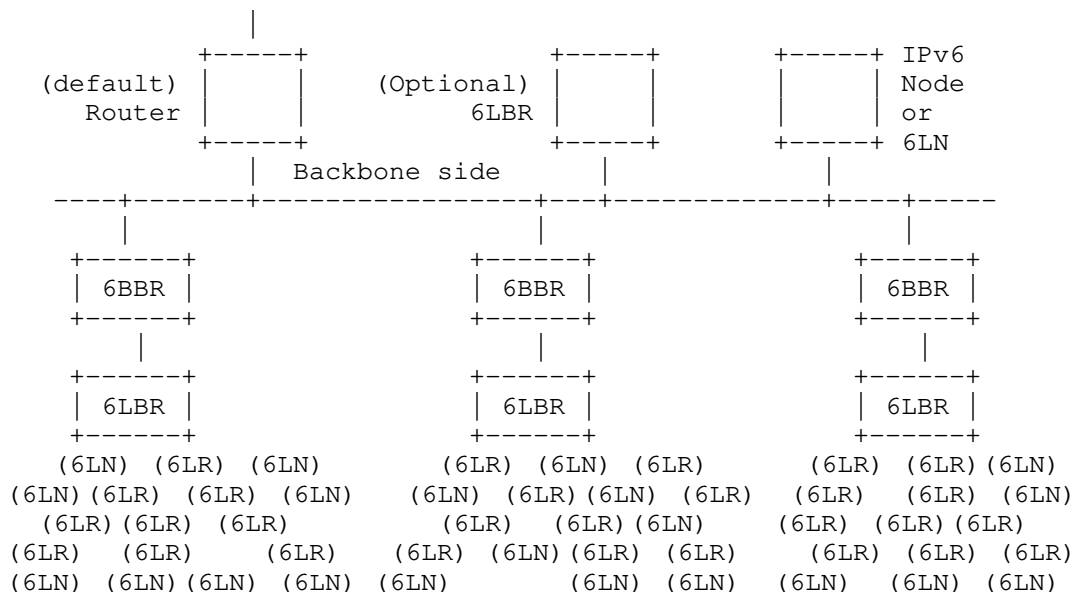


Figure 4: Route-Over Mesh Use case

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

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

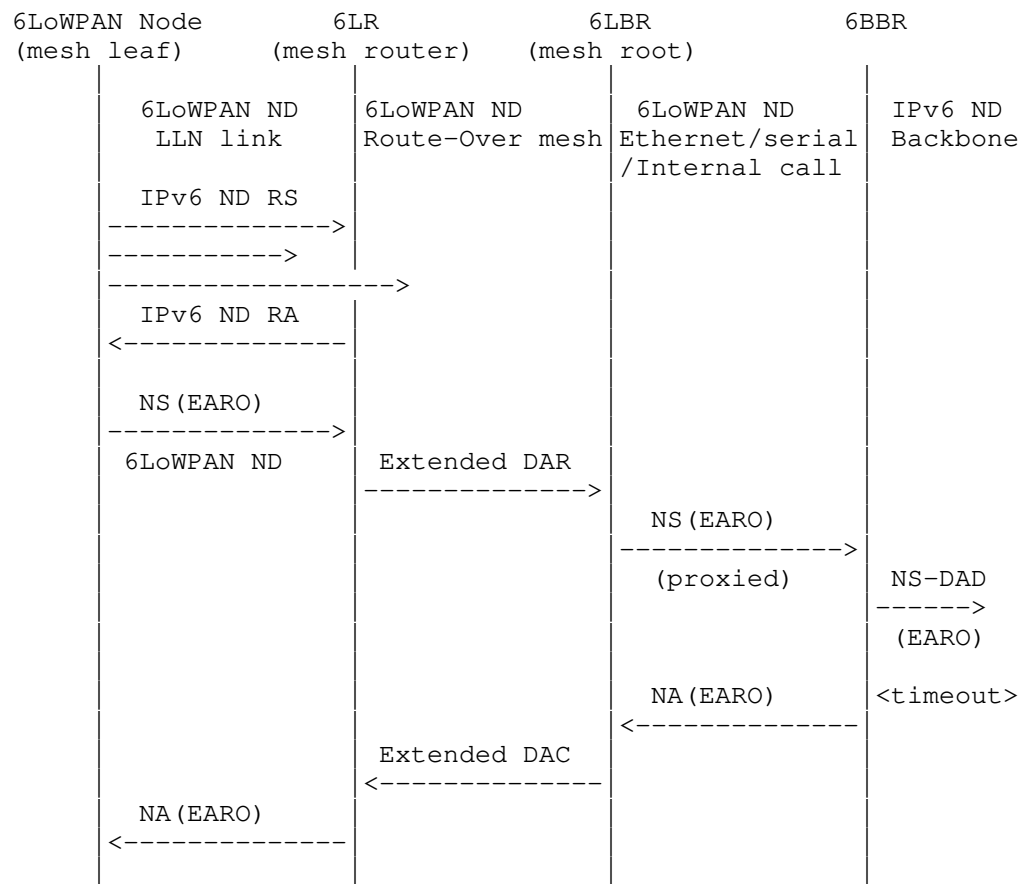


Figure 5: Initial Registration Flow over Route-Over Mesh

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

3.4. The Binding Table

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

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

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

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

Over the LLN, Binding Table management is as follows:

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

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

3.5. Primary and Secondary 6BBRs

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

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

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

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

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

3.6. Using Optimistic DAD

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

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

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

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

4. Multi-Link Subnet Considerations

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

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

5. Optional 6LBR serving the Multi-Link Subnet

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

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

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

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

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

6. Using IPv6 ND Over the Backbone Link

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

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

IPv6 ND operates as follows on the backbone:

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

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

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

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

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

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

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

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

In any state, the 6BBR acts as follows:

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

This behavior is specified in more detail in Section 9.

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

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

7. Routing Proxy Operations

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

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

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

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

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

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

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

8. Bridging Proxy Operations

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

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

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

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

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

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

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

9. Creating and Maintaining a Binding

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

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

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

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

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

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

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

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

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

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

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

9.1. Operations on a Binding in Tentative State

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

For a Binding in Tentative state:

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

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

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

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

9.2. Operations on a Binding in Reachable State

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

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

For a Binding in Reachable state:

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

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

9.3. Operations on a Binding in Stale State

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

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

For a Binding in Stale state:

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

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

10. Registering Node Considerations

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

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

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

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

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

11. Security Considerations

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

12. Protocol Constants

This Specification uses the following constants:

TENTATIVE_DURATION: 800 milliseconds

STALE_DURATION: see below

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

13. IANA Considerations

This document has no request to IANA.

14. Acknowledgments

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

15. Normative References

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

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

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

16. Informative References

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

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

[I-D.bi-savi-wlan]

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

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

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

[IEEEStd8021]

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

[IEEEStd80211]

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

[IEEEStd802151]

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

[IEEEStd802154]

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

Appendix A. Possible Future Extensions

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

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

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

Appendix B. Applicability and Requirements Served

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

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

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

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

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

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

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

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

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

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

Authors' Addresses

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

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

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

Email: charliep@computer.org

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

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

6lo
Internet-Draft
Updates: 6775 (if approved)
Intended status: Standards Track
Expires: December 23, 2017

P. Thubert, Ed.
cisco
E. Nordmark

S. Chakrabarti
June 21, 2017

An Update to 6LoWPAN ND
draft-ietf-6lo-rfc6775-update-06

Abstract

This specification updates RFC 6775 - 6LoWPAN Neighbor Discovery, to clarify the role of the protocol as a registration technique, simplify the registration operation in 6LoWPAN routers, as well as to provide enhancements to the registration capabilities and mobility detection for different network topologies including the backbone routers performing proxy Neighbor Discovery in a low power network.

Status of This Memo

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

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

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

This Internet-Draft will expire on December 23, 2017.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must

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

Table of Contents

1. Introduction	3
2. Applicability of Address Registration Options	3
3. Terminology	4
4. Updating RFC 6775	5
4.1. Extended Address Registration Option	6
4.2. Transaction ID	6
4.3. Owner Unique ID	7
4.4. Registering the Target Address	7
4.5. Link-Local Addresses and Registration	8
4.6. Maintaining the Registration States	9
5. Detecting Enhanced ARO Capability Support	10
6. Updated ND Options	11
6.1. The Enhanced Address Registration Option (EARO)	11
6.2. New 6LoWPAN capability Bits in the Capability Indication Option	14
7. Backward Compatibility	14
7.1. Discovering the capabilities of an ND peer	14
7.1.1. Using the E Flag in the CIO	14
7.1.2. Using the T Flag in the EARO	15
7.2. Legacy 6LoWPAN Node	15
7.3. Legacy 6LoWPAN Router	16
7.4. Legacy 6LoWPAN Border Router	16
8. Security Considerations	16
9. Privacy Considerations	18
10. IANA Considerations	18
11. Acknowledgments	20
12. References	20
12.1. Normative References	20
12.2. Informative References	21
12.3. External Informative References	23
Appendix A. Applicability and Requirements Served	24
Appendix B. Requirements	24
B.1. Requirements Related to Mobility	25
B.2. Requirements Related to Routing Protocols	25
B.3. Requirements Related to the Variety of Low-Power Link types	26
B.4. Requirements Related to Proxy Operations	27
B.5. Requirements Related to Security	27
B.6. Requirements Related to Scalability	28
Authors' Addresses	29

1. Introduction

The scope of this draft is an IPv6 Low Power Networks including star and mesh topologies. This specification modifies and extends the behavior and protocol elements of RFC 6775 "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)" [RFC6775] to enable additional capabilities such as:

- * Support the indication of mobility vs retry (T-bit)
- * Ease up requirement of registration for link-local addresses
- * Introducing Enhancement to Address Registration Option (ARO)
- * Permitting registration of target address
- * Clarification of support of privacy and temporary addresses

The following sections will discuss applicability of 6LoWPAN ND registration, new extensions and updates to RFC 6775. Finally, we will discuss how the extensions of registration framework can be useful for a scenario such as Backbone router(6BBR) proxy ND operations.

2. Applicability of Address Registration Options

The purpose of the Address Registration Option (ARO) [RFC6775] and of the Extended ARO (EARO) that is introduced in this document is to facilitate duplicate address detection (DAD) for hosts and pre-populate Neighbor Cache Entries (NCE) [RFC4861] in the routers to reduce the need for sending 'multicast neighbor solicitations' which may be harmful in low power constrained nodes networks where multicast is most often treated as broadcasts.

In some cases the address registration can fail or becomes useless for reasons other than a duplicate address. Examples are the router having run out of space, a registration bearing a stale sequence number (e.g. denoting a movement of the host after this registration was placed), a host misbehaving and attempting to register an invalid address such as the unspecified address [RFC4291], or the host using an address which is not topologically correct on that link. In such cases the host will receive an error to help diagnose the issue and may retry, possibly with a different address, and possibly registering to a different 6LR, depending on the returned error.

However, the ability to return errors to address registrations MUST NOT be used to restrict the ability of hosts to form and use addresses as recommended in "Host Address Availability

Recommendations" [RFC7934]. In particular, this is needed for enhanced privacy, which implies that each host will register a multiplicity of address as part mechanisms like "Privacy Extensions for Stateless Address Autoconfiguration (SLAAC) in IPv6" [RFC4941]. This implies that the capabilities of 6LR and 6LBRs in terms of number of registrations must be clearly announced in the router documentation, and that a network administrator should deploy adapted 6LR/6LBRs to support the number and type of devices in his network, based on the number of IPv6 addresses that those devices require.

3. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Readers are expected to be familiar with all the terms and concepts that are discussed in

"Neighbor Discovery for IP version 6" [RFC4861],

"IPv6 Stateless Address Autoconfiguration" [RFC4862],

"IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals" [RFC4919],

"Neighbor Discovery Optimization for Low-power and Lossy Networks" [RFC6775] and

"Multi-link Subnet Support in IPv6"
[I-D.ietf-ipv6-multilink-subnets].

as well as this additional terminology:

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

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

Extended LLN This is the aggregation of multiple LLNs as defined in RFC 4919 [RFC4919], interconnected by a Backbone Link via Backbone Routers, and forming a single IPv6 MultiLink Subnet.

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

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

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

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

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

4. Updating RFC 6775

This specification extends the Address Registration Option (ARO) defined in RFC 6775 [RFC6775]; in particular a "T" flag is added that must be set in NS messages when this specification is used, and echo'ed in NA messages to confirm that the protocol effectively supported. Support for this specification can thus be inferred from the presence of the Extended ARO ("T" flag set) in ND messages.

In order to support various types of link layers, this specification also adds recommendation to allow multiple registrations, including for privacy / temporary addresses, and provides new mechanisms to help clean up stale registration states as soon as possible.

A Registering Node that supports this specification will favor registering to a 6LR that indicates support for this specification over that of RFC 6775 [RFC6775].

4.1. Extended Address Registration Option

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

The address that is being registered with a Neighbor Solicitation (NS) with an EARO is now the Target Address, as opposed to the Source Address as specified in RFC 6775 [RFC6775] (see Section 4.4 for more). This change enables a 6LBR to use an address of his as source to the proxy-registration of an address that belongs to a LLN Node to a 6BBR. This also limits the use of an address as source address before it is registered and the associated Duplicate Address Detection (DAD) is complete.

The Unique ID in the EARO option does no more have to be a MAC address (see Section 4.3 for more). This enables in particular the use of a Provable Temporary UID (PT-UID) as opposed to burn-in MAC address, the PT-UID providing a trusted anchor by the 6LR and 6LBR to protect the state associated to the node.

The specification introduces a Transaction ID (TID) field in the EARO (see Section 4.2 for more on TID). The TID MUST be provided by a node that supports this specification and a new T flag MUST be set to indicate so. The T bit can be used to determine whether the peer supports this specification.

Finally, this specification introduces a number of new Status codes to help diagnose the cause of a registration failure (more in Table 1).

4.2. Transaction ID

The specification expects that the Registered Node can provide a sequence number called Transaction ID (TID) that is incremented with each re-registration. The TID is used to detect the freshness of the registration request and useful to detect one single registration by multiple 6LoWPAN border routers supporting the same large 6LoWPAN, as is the case for backbone routers (BBR).

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

Thus TID could be tracked to follow the sequence of mobility of a node. The details protocols of mobility verification by the border routers is not part of this specification.

4.3. Owner Unique ID

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

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

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

4.4. Registering the Target Address

This specification changes the behavior of the 6LN and the 6LR so that the Registered Address is found in the Target Address field of the NS and NA messages as opposed to the Source Address.

The reason for this change is to enable proxy-registrations on behalf of other nodes in Route-Over meshes, for instance to enable that a RPL root registers addresses on behalf LLN nodes that are deeper in a 6TiSCH mesh, as discussed in Appendix B.4. In that case, the Registering Node MUST indicate its own address as source of the ND message and its MAC address in the Source Link-Layer Address Option (SLLAO), since it still expects to get the packets and route them down the mesh. But the Registered Address belongs to another node, the Registered Node, and that address is indicated in the Target Address field of the NS message.

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

Since the Registering Node is the one that has reachability with the 6LR, and is the one expecting packets for the 6LN, it makes sense to maintain compatibility with RFC 6775 [RFC6775], and it is REQUIRED

that an SLLA Option is always placed in a registration NS(EARO) message.

4.5. Link-Local Addresses and Registration

Considering that LLN nodes are often not wired and may move, there is no guarantee that a Link-Local address stays unique between a potentially variable and unbounded set of neighboring nodes. Compared to RFC 6775 [RFC6775], this specification only requires that a Link-Local address is unique from the perspective of the peering nodes. This simplifies the Duplicate Address Detection (DAD) for Link-Local addresses, and there is no Duplicate Address Request (DAR) / Duplicate Address Confirmation (DAC) exchange between the 6LR and a 6LBR for Link-Local addresses.

Additionally, RFC 6775 [RFC6775] requires that a 6LoWPAN Node (6LN) uses an address being registered as the source of the registration message. This generates complexities in the 6LR to be able to cope with a potential duplication, in particular for global addresses. To simplify this, a 6LN and a 6LR that conform this specification always use Link-Local addresses as source and destination addresses for the registration NS/NA exchange. As a result, the registration is globally faster, and some of the complexity is removed.

In more details:

An exchange between two nodes using Link-Local addresses implies that they are reachable over one hop and that at least one of the 2 nodes acts as a 6LR. A node MUST register a Link-Local address to a 6LR in order to obtain reachability from that 6LR beyond the current exchange, and in particular to use the Link-Local address as source address to register other addresses, e.g. global addresses.

If there is no collision with an address previously registered to this 6LR by another 6LN, then, from the standpoint of this 6LR, this Link-Local address is unique and the registration is acceptable. Conversely, it may possibly happen that two different 6LRs expose a same Link-Local address but different link-layer addresses. In that case, a 6LN may only interact with one of the 6LR so as to avoid confusion in the 6LN neighbor cache.

The DAD process between the 6LR and a 6LoWPAN Border Router (6LBR), which is based on a Duplicate Address Request (DAR) / Duplicate Address Confirmation (DAC) exchange as described in RFC 6775 [RFC6775], does not need to take place for Link-Local addresses.

It is desired that a 6LR does not need to modify its state associated to the Source Address of an NS(EARO) message. For that reason, when

possible, it is RECOMMENDED to use an address that is already registered with a 6LR

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

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

Since there is no DAR/DAC exchange for Link-Local addresses, the 6LR may answer immediately to the registration of a Link-Local address, based solely on its existing state and the Source Link-Layer Option that MUST be placed in the NS(EARO) message as required in RFC 6775 [RFC6775].

A node needs to register its IPv6 Global Unicast IPv6 Addresses (GUA) to a 6LR in order to obtain a global reachability for these addresses via that 6LR. As opposed to a node that complies to RFC 6775 [RFC6775], a Registering Node registering a GUA does not use that GUA as Source Address for the registration to a 6LR that conforms this specification. The DAR/DAC exchange MUST take place for non-Link-Local addresses as prescribed by RFC 6775 [RFC6775].

4.6. Maintaining the Registration States

This section discusses protocol actions that involve the Registering Node, the 6LR and the 6LBR. It must be noted that the portion that deals with a 6LBR only applies to those addresses that are registered to it, which, as discussed in Section 4.5, is not the case for Link-Local addresses. The registration state includes all data that is stored in the router relative to that registration, in particular, but not limited to, an NCE in a 6LR. 6LBRs and 6BBRs may store additional registration information in more complex data structures and use protocols that are out of scope of this document to keep them synchronized when they are distributed.

When its Neighbor Cache is full, a 6LR cannot accept a new registration. In that situation, the EARO is returned in a NA message with a Status of 2, and the Registering Node may attempt to register to another 6LR. Conversely the registry in the 6LBR may be

saturated, in which case the 6LBR cannot guarantee that a new address is effectively not a duplicate. In that case, the 6LBR replies to a DAR message with a DAC message that carries a Status code 9 indicating "6LBR Registry saturated", and the address stays in TENTATIVE state.

A node renews an existing registration by repeatedly sending NS(EARO) messages for the Registered Address. In order to refresh the registration state in the 6LBR, these registrations MUST be reported to the 6LBR.

A node that ceases to use an address SHOULD attempt to deregister that address from all the 6LRs to which it has registered the address, which is achieved using an NS(EARO) message with a Registration Lifetime of 0.

A node that moves away from a particular 6LR SHOULD attempt to deregister all of its addresses registered to that 6LR.

Upon receiving a NS(EARO) message with a Registration Lifetime of 0 and determining that this EARO is the freshest for a given NCE (see Section 4.2), a 6LR cleans up its NCE. If the address was registered to the 6LBR, then the 6LR MUST report to the 6LBR, through a DAR/DAC exchange with the 6LBR, or an alternate protocol, indicating the null Registration Lifetime and the latest TID that this 6LR is aware of.

Upon the DAR message, the 6LBR evaluates if this is the freshest EARO it has received for that particular registry entry. If it is, then the entry is scheduled to be removed, and the DAR is answered with a DAC message bearing a Status of 0 "Success". If it is not the freshest, then a Status 2 "Moved" is returned instead, and the existing entry is conserved. The 6LBR SHOULD conserve the address in a DELAY state for a configurable period of time, so as to protect a mobile node that deregistered from one 6LR and did not register yet to a new one.

5. Detecting Enhanced ARO Capability Support

The nodes and routers in a network may be mixed and if a node wants to use EARO feature for address registration, it has to find a router which supports it. Thus all implementations with EARO option MUST provide the capability detection method using 6CIO option to support both types of registrations (ARO and EARO) as described in later sections. Moreover, any new implementation of 6LoWPAN is also RECOMMENDED to support 6LoWPAN Capability Indication option(6CIO) in general.

RFC 7400 [RFC7400] introduces the 6LoWPAN Capability Indication Option (6CIO) to indicate a node's capabilities to its peers. This specification extends the format defined in RFC 7400 to signal the support for EARO, as well as the capability to act as a 6LR, 6LBR and 6BBR.

With RFC 7400 [RFC7400], the 6CIO is typically sent Router Solicitation (RS) messages. When used to signal the capabilities above per this specification, the 6CIO is typically present Router Advertisement (RA) messages but can also be present in RS, Neighbor Solicitation (NS) and Neighbor Advertisement (NA) messages.

6. Updated ND Options

This specification does not introduce new options, but it modifies existing ones and updates the associated behaviors as follow:

6.1. The Enhanced Address Registration Option (EARO)

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

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

The EARO extends the ARO and is recognized by the "T" flag set.

When using the EARO option, the address being registered is found in the Target Address field of the NS and NA messages. This differs from 6LoWPAN ND RFC 6775 [RFC6775] which specifies that the address being registered is the source of the NS.

The format of the EARO option is as follows:

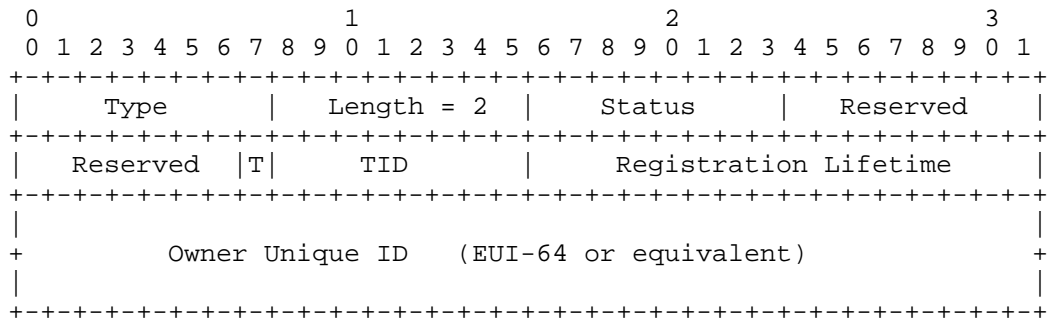


Figure 1: EARO

Option Fields

Type: 33

Length: 8-bit unsigned integer.

Status: 8-bit unsigned integer. Indicates the status of a registration in the NA response. MUST be set to 0 in NS messages. See Table 1 below.

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

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

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

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

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

Value	Description
0..2	See RFC 6775 [RFC6775]. Note that a Status of 1 "Duplicate Address" applies to the Registered Address. If the Source Address conflicts with an existing registration, "Duplicate Source Address" should be used.

3	Moved: The registration fails because it is not the freshest. This Status indicates that the registration is rejected because another more recent registration was done, as indicated by a same OUI and a more recent TID. One possible cause is a stale registration that has progressed slowly in the network and was passed by a more recent one. It could also indicate a OUI collision.
4	Removed: The binding state was removed. This may be placed in an asynchronous NS(ARO) message, or as the rejection of a proxy registration to a Backbone Router
5	Proof requested: The Registering Node is challenged for owning the Registered Address or for being an acceptable proxy for the registration. This Status is expected in asynchronous messages from a registrar (6LR, 6LBR, 6BBR) to indicate that the registration state is removed, for instance due to time out of a lifetime, or a movement. The receiver of the NA is the device that has performed a registration that is now stale and it should clean up its state.
6	Duplicate Source Address: The address used as source of the NS(ARO) conflicts with an existing registration.
7	Invalid Source Address: The address used as source of the NS(ARO) is not a Link-Local address as prescribed by this document.
8	Registered Address topologically incorrect: The address being registered is not usable on this link, e.g. it is not topologically correct
9	6LBR Registry saturated: A new registration cannot be accepted because the 6LBR Registry is saturated.
10	Incorrect proof: The proof of ownership of the registered address is not correct.

Table 1: EARO Status

Note: the code "6LBR Registry saturated" is used by 6LBRs instead of Status 2 when responding to a DAR/DAC exchange and passed on to the Registering Node by the 6LR. There is no point for the node to retry this registration immediately via another 6LR, since the problem is global to the network. The node may either abandon that address,

deregister other addresses first to make room, or keep the address in TENTATIVE state and retry later.

6.2. New 6LoWPAN capability Bits in the Capability Indication Option

This specification defines a number of capability bits in the CIO that was introduced by RFC 7400 [RFC7400].

Support for this specification is indicated by setting the "E" flag in a CIO option. Routers that are capable of acting as 6LR, 6LBR and 6BBR SHOULD set the L, B and P flags, respectively.

Those flags are not mutually exclusive and if a router is capable of multiple roles, it SHOULD set all the related flags.




Figure 2: New capability Bits L, B, P, E in the CIO

Option Fields

Type: 36

L: Node is a 6LR, it can take registrations.

B: Node is a 6LBR.

P: Node is a 6BBR, proxying for nodes on this link.

E: This specification is supported and applied.

7. Backward Compatibility

7.1. Discovering the capabilities of an ND peer

7.1.1. Using the E Flag in the CIO

If the CIO is used in an ND message, then the "E" Flag MUST be set by the sending node if supports this specification.

It is RECOMMENDED that a router that supports this specification indicates so with a CIO option, but this might not be practical if the link-layer MTU is too small.

If the Registering Node receives a CIO in a RA, then the setting of the E" Flag indicates whether or not this specification is supported.

A node which does not implement this draft or parse 6CIO option, MUST ignore the packet and the sender of option SHOULD use legacy registration method according to RFC 6775 [RFC6775] after a timeout period.

7.1.2. Using the T Flag in the EARO

One alternate way for a 6LN to discover the router's capabilities to first register a Link Local address, placing the same address in the Source and Target Address fields of the NS message, and setting the "T" Flag. The node may for instance register an address that is based on EUI-64. For such address, DAD is not required and using the SLLAO option in the NS is actually more amenable with existing ND specifications such as the "Optimistic Duplicate Address Detection (DAD) for IPv6" [RFC4429]. Once that first registration is complete, the node knows from the setting of the "T" Flag in the response whether the router supports this specification. If this is verified, the node may register other addresses that it owns, or proxy-register addresses on behalf some another node, indicating those addresses being registered in the Target Address field of the NS messages, while using one of its own, already registered, addresses as source.

A node that supports this specification MUST always use an EARO as a replacement to an ARO in its registration to a router. This is harmless since the "T" flag and TID field are reserved in RFC 6775 [RFC6775] are ignored by a legacy router. A router that supports this specification answers to an ARO with an ARO and to an EARO with an EARO.

This specification changes the behavior of the peers in a registration flows. To enable backward compatibility, a node that registers to a router that is not known to support this specification MUST behave as prescribed by RFC 6775. Once the router is known to support this specification, the node MUST obey this specification.

7.2. Legacy 6LoWPAN Node

A legacy 6LN will use the Registered Address as source and will not use an EARO option. In order to be backward compatible, an updated 6LR needs to accept that registration if it is valid per the RFC 6775 [RFC6775] specification, and manage the binding cache accordingly.

The main difference with RFC 6775 is that DAR/DAC exchange for DAD may be avoided for Link-Local addresses. Additionally, the 6LR SHOULD use an EARO in the reply, and may use any of the Status codes defined in this specification.

7.3. Legacy 6LoWPAN Router

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

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

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

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

7.4. Legacy 6LoWPAN Border Router

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

8. Security Considerations

This specification extends RFC 6775 [RFC6775], and the security section of that draft also applies to this as well. In particular, it is expected that the link layer is sufficiently protected to prevent a rogue access, either by means of physical or IP security on the Backbone Link and link layer cryptography on the LLN. This specification also expects that the LLN MAC provides secure unicast to/from the Backbone Router and secure Broadcast from the Backbone Router in a way that prevents tempering with or replaying the RA messages.

This specification recommends to using privacy techniques (more in section Section 9, and protection against address theft such as provided by "Address Protected Neighbor Discovery for Low-power and

Lossy Networks" [I-D.ietf-6lo-ap-nd], which guarantees the ownership of the Registered Address using a cryptographic OUID.

The registration mechanism may be used by a rogue node to attack the 6LR or the 6LBR with a Denial-of-Service attack against the registry. It may also happen that the registry of a 6LR or a 6LBR is saturated and cannot take any more registration, which effectively denies the requesting a node the capability to use a new address. In order to alleviate those concerns, Section 4.6 provides a number of recommendations that ensure that a stale registration is removed as soon as possible from the 6LR and 6LBR. In particular, this specification recommends that:

- o A node that ceases to use an address should attempt to deregister that address from all the 6LRs to which it is registered. The flow is propagated to the 6LBR when needed, and a sequence number is used to make sure that only the freshest command is acted upon.
- o The nodes should be configured with a Registration Lifetime that reflects their expectation of how long they will use the address with the 6LR to which it is registered. In particular, use cases that involve mobility or rapid address changes should use lifetimes that are homogeneous with the expectation of presence.
- o The router (6LR or 6LBR) should be configurable so as to limit the number of addresses that can be registered by a single node, as identified at least by MAC address and preferably by security credentials. When that maximum is reached, the router should use a Least-Recently-Used (LRU) logic so as to clean up the addresses that were not used for the longest time, keeping at least one Link-Local address, and attempting to keep one or more stable addresses if such can be recognized, e.g. from the way the IID is formed or because they are used over a much longer time span than other (privacy, shorter-lived) addresses.
- o Administrators should take great care to deploy adequate numbers of 6LR to cover the needs of the nodes in their range, so as to avoid a situation of starving nodes. It is expected that the 6LBR that serves a LLN is a more capable node than the average 6LR, but in a network condition where it may become saturated, a particular deployment should distribute the 6LBR functionality, for instance by leveraging a high speed Backbone and Backbone Routers to aggregate multiple LLNs into a larger subnet.

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

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

9. Privacy Considerations

As indicated in section Section 2, this protocol does not aim at limiting the number of IPv6 addresses that a device can form. A host should be able to form and register any address that is topologically correct in the subnet(s) advertised by the 6LR/6LBR.

This specification does not mandate any particular way for forming IPv6 addresses, but it recognizes that use of EUI-64 for forming the Interface ID in the Link-Local address prevents the usage of "SEcure Neighbor Discovery (SEND)" [RFC3971] and "Cryptographically Generated Addresses (CGA)" [RFC3972], and that of address privacy techniques.

"Privacy Considerations for IPv6 Adaptation-Layer Mechanisms" [RFC8065] addresses why privacy is important and how to form such addresses. All implementations and deployment must consider the option of privacy addresses in their own environment. Also future specifications involving 6LoWPAN Neighbor Discovery should consult "Recommendation on Stable IPv6 Interface Identifiers" [RFC8064] for default interface identification.

10. IANA Considerations

IANA is requested to create a new subregistry for "ARO Flags" under the "Internet Control Message Protocol version 6 (ICMPv6) Parameters". This specification defines 8 positions, bit 0 to bit 7, and assigns bit 7 for the "T" flag in Section 6.1. The policy is "IETF Review" or "IESG Approval" [RFC5226]. The initial content of the registry is as shown in Table 2.

New subregistry for ARO Flags under the "Internet Control Message Protocol version 6 (ICMPv6) Parameters"

ARO Status	Description	Document
0..6	Unassigned	
7	"T" Flag	RFC This

Table 2: new ARO Flags

IANA is requested to make additions to existing registries as follows:

Address Registration Option Status Values Registry

ARO Status	Description	Document
3	Moved	RFC This
4	Removed	RFC This
5	Proof requested	RFC This
6	Duplicate Source Address	RFC This
7	Invalid Source Address	RFC This
8	Registered Address topologically incorrect	RFC This
9	6LBR registry saturated	RFC This
10	Incorrect proof	RFC This

Table 3: New ARO Status values

Subregistry for "6LoWPAN capability Bits" under the "Internet Control Message Protocol version 6 (ICMPv6) Parameters"

capability Bit	Description	Document
11	6LR capable (L bit)	RFC This
12	6LBR capable (B bit)	RFC This
13	6BBR capable (P bit)	RFC This
14	EARO support (E bit)	RFC This

Table 4: New 6LoWPAN capability Bits

11. Acknowledgments

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

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<http://www.rfc-editor.org/info/rfc4291>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<http://www.rfc-editor.org/info/rfc4861>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <<http://www.rfc-editor.org/info/rfc4862>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 5226, DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<http://www.rfc-editor.org/info/rfc6282>>.
- [RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <<http://www.rfc-editor.org/info/rfc6775>>.
- [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, <<http://www.rfc-editor.org/info/rfc7400>>.

12.2. Informative References

- [I-D.chakrabarti-nordmark-6man-efficient-nd]
Chakrabarti, S., Nordmark, E., Thubert, P., and M. Wasserman, "IPv6 Neighbor Discovery Optimizations for Wired and Wireless Networks", draft-chakrabarti-nordmark-6man-efficient-nd-07 (work in progress), February 2015.
- [I-D.delcarpio-6lo-wlanah]
Vega, L., Robles, I., and R. Morabito, "IPv6 over 802.11ah", draft-delcarpio-6lo-wlanah-01 (work in progress), October 2015.
- [I-D.ietf-6lo-ap-nd]
Sarikaya, B., Thubert, P., and M. Sethi, "Address Protected Neighbor Discovery for Low-power and Lossy Networks", draft-ietf-6lo-ap-nd-02 (work in progress), May 2017.
- [I-D.ietf-6lo-backbone-router]
Thubert, P., "IPv6 Backbone Router", draft-ietf-6lo-backbone-router-03 (work in progress), January 2017.
- [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-07 (work in progress), June 2017.
- [I-D.ietf-6tisch-architecture]
Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", draft-ietf-6tisch-architecture-11 (work in progress), January 2017.
- [I-D.ietf-bier-architecture]
Wijnands, I., Rosen, E., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast using Bit Index Explicit Replication", draft-ietf-bier-architecture-07 (work in progress), June 2017.
- [I-D.ietf-ipv6-multilink-subnets]
Thaler, D. and C. Huitema, "Multi-link Subnet Support in IPv6", draft-ietf-ipv6-multilink-subnets-00 (work in progress), July 2002.

- [I-D.popa-6lo-6loplc-ipv6-over-ieee19012-networks]
Popa, D. and J. Hui, "6LoPLC: Transmission of IPv6 Packets over IEEE 1901.2 Narrowband Powerline Communication Networks", draft-popa-6lo-6loplc-ipv6-over-ieee19012-networks-00 (work in progress), March 2014.
- [RFC3610] Whiting, D., Housley, R., and N. Ferguson, "Counter with CBC-MAC (CCM)", RFC 3610, DOI 10.17487/RFC3610, September 2003, <<http://www.rfc-editor.org/info/rfc3610>>.
- [RFC3810] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, DOI 10.17487/RFC3810, June 2004, <<http://www.rfc-editor.org/info/rfc3810>>.
- [RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "Secure Neighbor Discovery (SEND)", RFC 3971, DOI 10.17487/RFC3971, March 2005, <<http://www.rfc-editor.org/info/rfc3971>>.
- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, DOI 10.17487/RFC3972, March 2005, <<http://www.rfc-editor.org/info/rfc3972>>.
- [RFC4429] Moore, N., "Optimistic Duplicate Address Detection (DAD) for IPv6", RFC 4429, DOI 10.17487/RFC4429, April 2006, <<http://www.rfc-editor.org/info/rfc4429>>.
- [RFC4919] Kushalnagar, N., Montenegro, G., and C. Schumacher, "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals", RFC 4919, DOI 10.17487/RFC4919, August 2007, <<http://www.rfc-editor.org/info/rfc4919>>.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, DOI 10.17487/RFC4941, September 2007, <<http://www.rfc-editor.org/info/rfc4941>>.
- [RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <<http://www.rfc-editor.org/info/rfc6550>>.

- [RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <<http://www.rfc-editor.org/info/rfc7217>>.
- [RFC7428] Brandt, A. and J. Buron, "Transmission of IPv6 Packets over ITU-T G.9959 Networks", RFC 7428, DOI 10.17487/RFC7428, February 2015, <<http://www.rfc-editor.org/info/rfc7428>>.
- [RFC7668] Nieminen, J., Savolainen, T., Isomaki, M., Patil, B., Shelby, Z., and C. Gomez, "IPv6 over BLUETOOTH(R) Low Energy", RFC 7668, DOI 10.17487/RFC7668, October 2015, <<http://www.rfc-editor.org/info/rfc7668>>.
- [RFC7934] Colitti, L., Cerf, V., Cheshire, S., and D. Schinazi, "Host Address Availability Recommendations", BCP 204, RFC 7934, DOI 10.17487/RFC7934, July 2016, <<http://www.rfc-editor.org/info/rfc7934>>.
- [RFC8064] Gont, F., Cooper, A., Thaler, D., and W. Liu, "Recommendation on Stable IPv6 Interface Identifiers", RFC 8064, DOI 10.17487/RFC8064, February 2017, <<http://www.rfc-editor.org/info/rfc8064>>.
- [RFC8065] Thaler, D., "Privacy Considerations for IPv6 Adaptation-Layer Mechanisms", RFC 8065, DOI 10.17487/RFC8065, February 2017, <<http://www.rfc-editor.org/info/rfc8065>>.
- [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, <<http://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, <<http://www.rfc-editor.org/info/rfc8163>>.

12.3. External Informative References

- [IEEEstd802154]
IEEE, "IEEE Standard for Low-Rate Wireless Networks", IEEE Standard 802.15.4, DOI 10.1109/IEEESTD.2016.7460875, <<http://ieeexplore.ieee.org/document/7460875/>>.

Appendix A. Applicability and Requirements Served

This specification extends 6LoWPAN ND to sequence the registration and serves the requirements expressed Appendix B.1 by enabling the mobility of devices from one LLN to the next based on the complementary work in the "IPv6 Backbone Router" [I-D.ietf-6lo-backbone-router] specification.

In the context of the the TimeSlotted Channel Hopping (TSCH) mode of IEEE Std. 802.15.4 [IEEEstd802154], the "6TiSCH architecture" [I-D.ietf-6tisch-architecture] introduces how a 6LoWPAN ND host could connect to the Internet via a RPL mesh Network, but this requires additions to the 6LoWPAN ND protocol to support mobility and reachability in a secured and manageable environment. This specification details the new operations that are required to implement the 6TiSCH architecture and serves the requirements listed in Appendix B.2.

The term LLN is used loosely in this specification to cover multiple types of WLANs and WPANs, including Low-Power Wi-Fi, BLUETOOTH(R) Low Energy, IEEE Std.802.11AH and IEEE Std.802.15.4 wireless meshes, so as to address the requirements discussed in Appendix B.3

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

"Efficiency aware IPv6 Neighbor Discovery Optimizations" [I-D.chakrabarti-nordmark-6man-efficient-nd] suggests that 6LoWPAN ND [RFC6775] can be extended to other types of links beyond IEEE Std. 802.15.4 for which it was defined. The registration technique is beneficial when the Link-Layer technique used to carry IPv6 multicast packets is not sufficiently efficient in terms of delivery ratio or energy consumption in the end devices, in particular to enable energy-constrained sleeping nodes. The value of such extension is especially apparent in the case of mobile wireless nodes, to reduce the multicast operations that are related to classical ND ([RFC4861], [RFC4862]) and plague the wireless medium. This serves scalability requirements listed in Appendix B.6.

Appendix B. Requirements

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

B.1. Requirements Related to Mobility

Due to the unstable nature of LLN links, even in a LLN of immobile nodes a 6LN may change its point of attachment to a 6LR, say 6LR-a, and may not be able to notify 6LR-a. Consequently, 6LR-a may still attract traffic that it cannot deliver any more. When links to a 6LR change state, there is thus a need to identify stale states in a 6LR and restore reachability in a timely fashion.

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

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

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

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

B.2. Requirements Related to Routing Protocols

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

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

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

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

Related requirements are:

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

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

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

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

6LoWPAN ND [RFC6775] was defined with a focus on IEEE Std.802.15.4 and in particular the capability to derive a unique Identifier from a globally unique MAC-64 address. At this point, the 6lo Working Group is extending the 6LoWPAN Header Compression (HC) [RFC6282] technique to other link types ITU-T G.9959 [RFC7428], Master-Slave/Token-Passing [RFC8163], DECT Ultra Low Energy [RFC8105], Near Field Communication [I-D.ietf-6lo-nfc], IEEE Std. 802.11ah [I-D.delcarpio-6lo-wlanah], as well as IEEE1901.2 Narrowband Powerline Communication Networks [I-D.popa-6lo-6loplc-ipv6-over-ieee19012-networks] and BLUETOOTH(R) Low Energy [RFC7668].

Related requirements are:

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

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

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

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

B.4. Requirements Related to Proxy Operations

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

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

Related requirements are:

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

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

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

B.5. Requirements Related to Security

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

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

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

Related requirements are:

Req5.1: 6LoWPAN ND security mechanisms SHOULD provide a mechanism for the 6LR, 6LBR and 6BBR to authenticate and authorize one another for their respective roles, as well as with the 6LoWPAN Node for the role of 6LR.

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

Req5.3: 6LoWPAN ND security mechanisms SHOULD lead to small packet sizes. In particular, the NS, NA, DAR and DAC messages for a re-registration flow SHOULD NOT exceed 80 octets so as to fit in a secured IEEE Std.802.15.4 [IEEEstd802154] frame.

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

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

Req5.6: The 6LoWPAN ND security mechanisms SHOULD enable the variation of CCM [RFC3610] called CCM* for use at both Layer 2 and Layer 3, and SHOULD enable the reuse of security code that has to be present on the device for upper layer security such as TLS.

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

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

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

B.6. Requirements Related to Scalability

Use cases from Automatic Meter Reading (AMR, collection tree operations) and Advanced Metering Infrastructure (AMI, bi-directional communication to the meters) indicate the needs for a large number of

LLN nodes pertaining to a single RPL DODAG (e.g. 5000) and connected to the 6LBR over a large number of LLN hops (e.g. 15).

Related requirements are:

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

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

Authors' Addresses

Pascal Thubert (editor)
Cisco Systems, Inc
Sophia Antipolis
FRANCE

Email: pthubert@cisco.com

Erik Nordmark
Santa Clara, CA
USA

Email: nordmark@sonic.net

Samita Chakrabarti
San Jose, CA
USA

Email: samitac.ietf@gmail.com

6Lo Working Group
Internet-Draft
Intended status: Informational
Expires: January 4, 2018

Y-G. Hong
ETRI
C. Gomez
UPC/i2cat
Y-H. Choi
ETRI
D-Y. Ko
SKtelecom
AR. Sangi
Individual Contributor
T. Aanstoot
Modio AB
S. Chakrabarti
July 3, 2017

IPv6 over Constrained Node Networks (6lo) Applicability & Use cases
draft-ietf-6lo-use-cases-02

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), 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 link layer networks connecting devices to each other or to each cloud.

Status of This Memo

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

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

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

This Internet-Draft will expire on January 4, 2018.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Conventions and Terminology	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)	6
3.7. IEEE 802.15.4e (specified)	7
3.8. LTE MTC (example of a potential candidate)	8
3.9. 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
5. Design Space and Guidelines for 6lo Deployment	12
5.1. Design Space Dimensions for 6lo Deployment	12
5.2. Guidelines for adopting IPv6 stack (6lo/6LoWPAN)	14
6. 6lo Use Case Examples	16
7. IANA Considerations	16
8. Security Considerations	17
9. Acknowledgements	17
10. References	17
10.1. Normative References	17
10.2. Informative References	19
Appendix A. Other 6lo Use Case Examples	21
A.1. Use case of ITU-T G.9959: Smart Home	21
A.2. Use case of DECT-ULE: Smart Home	22
A.3. Use case of MS/TP: Management of District Heating	22
A.4. Use case of NFC: Alternative Secure Transfer	23
A.5. Use case of PLC: Smart Grid	23

A.6. Use case of IEEE 802.15.4e: Industrial Automation	24
Authors' Addresses	25

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], 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 is new to IPv6-over-lowpower networks concept and wants to assess if variance of 6LoWPAN stack [6lo] can be applied to the constrained 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. And it described 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). 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.

Devices such as mobile phones, notebooks, tablets and other handheld computing devices which will include Bluetooth 4.1 chipsets will probably also have the low-energy variant of Bluetooth. Bluetooth LE will also be 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.

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)

MS/TP is a contention-free access method for the RS-485 physical layer, which is used extensively in building automation networks.

An MS/TP device is typically based on a low-cost microcontroller with limited processing power and memory. Together with low data rates and a small address space, these constraints are similar to those faced in 6lowpan networks and suggest some elements of that solution might be leveraged. MS/TP differs significantly from 6lowpan in at least three aspects: a) MS/TP devices typically have a continuous source of power, b) all MS/TP devices on a segment can communicate directly so there are no hidden node or mesh routing issues, and c)

recent changes to MS/TP provide support for large payloads, eliminating the need for link-layer fragmentation and reassembly.

MS/TP is designed to enable multidrop networks over shielded twisted pair wiring, although not according to standards, in lower speeds, normally 9600 bit/s, re-purposed telecom wiring is widely in use, keeping deployment cost down. It can support a data rate of 115,200 baud on segments up to 1000 meters in length, or segments up to 1200 meters in length at lower baud rates. An MS/TP link requires only a UART, an RS-485 transceiver with a driver that can be disabled, and a 5ms resolution timer. These features make MS/TP a cost-effective and very reliable field bus for the most numerous and least expensive devices in a building automation network [RFC8163]. MS/TP can be used for the management of district heating.

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)

Unlike other dedicated communication infrastructure, the required medium (power conductor) is 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.

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 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 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-transmitt 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. LTE MTC (example of a potential candidate)

LTE category defines the overall performance and capabilities of the UE(User Equipment). For example, the maximum down rate of category 1 UE and category 2 UE are 10.3 Mbit/s and 51.0 Mbit/s respectively. There are many categories in LTE standard. 3GPP standards defined the category 0 to be used for low rate IoT service in release 12. Since category 1 and category 0 could be used for low rate IoT service, these categories are called LTE MTC (Machine Type Communication) [LTE_MTC].

LTE MTC offer advantages in comparison to above category 2 and is appropriate to be used for low rate IoT services such as low power and low cost. LTE MTC can be used for a gateway of a wireless bachhaul network.

3.9. 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 paramters 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	District Heating	Health-care Service	Smart Grid	Industrial Automation
Topology & Subnet	L2-mesh or L3-mesh	Star No mesh	Star No mesh	Bus MS/TP	P2P L2-mesh	Tree No mesh	Mesh
Mobility Reqmt	No	Low	No	No	Moderate	No	No
Security Reqmt	High + Privacy required	Partially	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	Infrequent	Infrequent	Infrequent	Frequent	Small	Infrequent	Infrequent
RFC # or Draft	RFC7428	RFC7668	RFC8105	RFC8163	6lo-nfc	hou-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, ICPMv6 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
- o 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.

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: Originally, the link layer technologies of 6lo have low rate of data transmission. But, by adjusting the MTU, it can deliver higher 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 is dependent on the 6lo use case. If the 6lo nodes can move or moved around, it requires a mobility management mechanism.
- 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 [I-D.ietf-lwig-energy-efficient]. 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.

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

The following guideline targets candidates for new constrained L2 technologies that consider running modified 6LoWPAN stack. 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 lowpower 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 of a packet (MTU) 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 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 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 quite 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.

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 (Ministerio de Educacion, Cultura y Deporte) through the Jose Castillejo grant CAS15/00336. 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.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC4919] Kushalnagar, N., Montenegro, G., and C. Schumacher, "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals", RFC 4919, DOI 10.17487/RFC4919, August 2007, <<http://www.rfc-editor.org/info/rfc4919>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<http://www.rfc-editor.org/info/rfc4944>>.
- [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, <<http://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, <<http://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, <<http://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, <<http://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, <<http://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, <<http://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, <<http://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, <<http://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, <<http://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, <<http://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, <<http://www.rfc-editor.org/info/rfc8036>>.
- [RFC8065] Thaler, D., "Privacy Considerations for IPv6 Adaptation-Layer Mechanisms", RFC 8065, DOI 10.17487/RFC8065, February 2017, <<http://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, <<http://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, <<http://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, <<http://www.rfc-editor.org/info/rfc8163>>.

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, <<http://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, <<http://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-07 (work in progress),
June 2017.
- [I-D.ietf-lwig-energy-efficient]
Gomez, C., Kovatsch, M., Tian, H., and Z. Cao, "Energy-
Efficient Features of Internet of Things Protocols",
draft-ietf-lwig-energy-efficient-07 (work in progress),
March 2017.
- [I-D.ietf-roll-aodv-rpl]
Anamalamudi, S., Zhang, M., Sangi, A., Perkins, C., and S.
Anand, "Asymmetric AODV-P2P-RPL in Low-Power and Lossy
Networks (LLNs)", draft-ietf-roll-aodv-rpl-01 (work in
progress), March 2017.
- [I-D.ietf-6tisch-6top-sf0]
Dujovne, D., Grieco, L., Palattella, M., and N. Accettura,
"6TiSCH 6top Scheduling Function Zero (SF0)", draft-ietf-
6tisch-6top-sf0-04 (work in progress), April 2017.
- [I-D.satish-6tisch-6top-sf1]
Anamalamudi, S., Zhang, M., Sangi, A., Perkins, C., and S.
Anand, "Scheduling Function One (SF1) for hop-by-hop
Scheduling in 6tisch Networks", draft-satish-6tisch-6top-
sf1-03 (work in progress), February 2017.
- [IETF_6lo]
"IETF IPv6 over Networks of Resource-constrained Nodes
(6lo) working group",
<<https://datatracker.ietf.org/wg/6lo/charter/>>.
- [G.9959]
"International Telecommunication Union, "Short range
narrow-band digital radiocommunication transceivers - PHY
and MAC layer specifications", ITU-T Recommendation",
January 2015.
- [LTE_MTC]
"3GPP TS 36.306 V13.0.0, 3rd Generation Partnership
Project; Technical Specification Group Radio Access
Network; Evolved Universal Terrestrial Radio Access
(E-UTRA); User Equipment (UE) radio access capabilities
(Release 13)", December 2015.

[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>>.

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: Management of District Heating

The key feature of MS/TP is its ability to run on the same cabling as BACnet and some use of ModBus, the defacto standard for low bandwidth industry communication. Specially Modbus has been around since the 1980 and is still the standard for talking to fans, heat pumps, water purifying equipment and everything else delivering electricity, clean water and ventilation.

Example: Use of MS/TP for management of district heating

The mechanical room in the cellar of an apartment building gets district heating and electricity from the utility providers. The room has a Supervisory Control And Data Acquisition (SCADA) computer talking to a centralized server and command center somewhere else over IP, on the other hand it is controlling the heating, fans and distribution panel over a 2-wire RS-485 based protocol to make sure the logic controller for district heating keeps a constant temperature at the tapwater, the logic controller for heat production

keeps the right radiator temperature depending on the weather and the fans have a correct speed and are switched off in case district heating fails to prevent cooling out the building and give certain commands in case smoke is detected. Speed is not important, in this usecase, 19,200 bit/s capable equipment is sold as high speed communication capable. Reliability is important, this not working will easily give millions of dollars of damage. Normally the setup is that the SCADA device asks a question to a specific controlling device, gets an answer from the controlling device, asks a new question to some other device.

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-sf0] 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

Deoknyong Ko
SKtelecom
9-1 Byundang-gu Sunae-dong, Seongnam-si
Gyeonggi-do 13595
Korea

Phone: +82 10 3356 8052
Email: engineer@sk.com

Abdur Rashid Sangi
Individual Contributor

Email: sangi_bahrian@yahoo.com

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

Email: take@modio.se

Samita Chakrabarti
San Jose, CA
USA

Email: samitac.ietf@gmail.com

6Lo Working Group
Internet-Draft
Intended status: Informational
Expires: December 15, 2017

MS. Akbar
Bournemouth University
R. Bin Rais
Ajman University
AR. Sangi
Individual Contributor
M. Zhang
Huawei Technologies
C. Perkins
Futurewei
June 13, 2017

Transmission of IPv6 Packets over Wireless Body Area Networks (WBANs)
draft-sajjad-6lo-wban-00

Abstract

Wireless Body Area Networks (WBANs) intend to facilitate use cases related to medical field. IEEE 802.15.6 defines PHY and MAC layer and is designed to deal with better penetration through the human tissue without creating any damage to human tissues with the approved MICS (Medical Implant Communication Service) band by USA Federal Communications Commission (FCC). Devices in WBANs conform to this IEEE standard.

This specification defines details to enable transmission of IPv6 packets, method of forming link-local and statelessly autoconfigured IPv6 addresses on WBANs.

Status of This Memo

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

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

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

This Internet-Draft will expire on December 15, 2017.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Conventions and Terminology	3
3. Use cases for IEEE 802.15.6	3
3.1. Hospital Patient Monitoring	3
3.2. Patient monitoring for Chronic Diseases	5
3.3. Elderly Patient Monitoring	5
4. Why 6lo is required for IEEE 802.15.6	5
4.1. IPv6 Connectivity requirements	5
4.2. Limited Packet Size	6
4.3. Topology requirements	6
5. Scope/Purpose	6
6. Layer 2 Overview	7
6.1. Frame format	8
6.2. Frequency bands	8
6.3. Channel modes of IEEE 802.15.6	10
7. --IETF to standardize--	11
8. IANA Considerations	11
9. Security Considerations	12
10. Acknowledgements	12
11. References	12
11.1. Normative References	12
11.2. Informative References	12
Appendix A. Patient monitoring use case - Spoke Hub	13
Appendix B. Patient monitoring use case - Connected	15
Authors' Addresses	16

1. Introduction

Wireless Body Area Networks (WBANs) are comprised of devices that conform to the [IEEE802.15.6], standard by the IEEE. IEEE 802.15.6 provides specification for the MAC layer to access the channel. The

coordinator divides the channel into superframe time structures to allocate resources [SURVEY-WBAN] [MAC-WBAN]. Superframes are bounded by equal length beacons through the coordinator. Usually beacons are sent at beacon periods except inactive superframes or limited by regulation. This standard works under following three channel access modes.

Task group for 802.15.6 was established by IEEE in November 2007 for standardisation of WBANs and it was approved in 2012. This standard works in and around human body and focus on operating at lower frequencies and short range. The focus of this standard is to design a communication standard for MAC and physical layer to support different applications, namely, medical and no-medical applications. Medical applications refer to collection of vital information in real time (monitoring) for diagnoses and treatment of various diseases with help of different sensors (accelerometer, temperature, BP and EMG etc.). It defines a MAC layer that can operate with three different PHY layers i.e. human body communication (HBC), ultra-wideband (UWB) and Narrowband (NB). IEEE 802.15.6 provides specification for MAC layer to access the channel. The coordinator divides the channel into superframe time structures to allocate resources. Superframes are bounded by equal length beacons through coordinator. The purpose of the draft is to highlight the need of IEEE 802.15.6 for WBASNs and its integration issues while connecting it with IPv6 network. The use cases are provided to elaborate the scenarios with implantable and wearable biomedical sensors. 6lowpan provides IPv6 connectivity for IEEE 802.15.4; however, it will not work with IEEE 802.15.6 due to the difference in frame format in terms of size and composition.

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. Use cases for IEEE 802.15.6

3.1. Hospital Patient Monitoring

In the hospital environment, several levels of patient monitoring services are required as different patients needs different monitoring services e.g., a patient in Intensive Care Unit (ICU) requires high prioritized periodic data services with limited delay and high throughput than the patient in a normal ward. Usually, a patient is equipped with multiple sensors that measure vital signals like heart activity, muscle movements, blood pressure, body oxygen level and brain stimulation via integrated sensors i.e.,

(Electrocardiography), BP (Blood Pressure) monitor, EMG (Electromyography), pulse oximeter and EEG (Electroencephalography) etc. These sensors are categorized as wearable and implantable sensors, hence we are assuming that equipped sensors are mixture of wearable and implantable sensors which creates restriction to use IEEE 802.15.6 as it is designed to deal with better penetration through the human tissue without creating any damage to human tissues with the approved MICS band by USA Federal Communications Commission (FCC). In a hospital use case scenario, the initial data generated by numerous biomedical sensor nodes is collected by a central coordinator.

In this case, Table 3 presents the summary of traffic patterns for different biomedical sensor nodes attached to human body with data generation rate, required data rate from channel and QoS requirements.

Sensor Nodes	Data Generation Interval	Required Data Rate (Kbps)	Delay Requirement	Power Consumption	Reliability Requirement
ECG	4 ms	34	<125ms	Low	High
EMG	6 ms	19.6	<125ms	Low	High
EEG	4 ms	19.6	<125ms	Low	High
SpO2 (Pulse Oximeter)	10 ms	13.2	<250ms	Low	Medium
BP	10 ms	13.2	<250ms	Medium	Medium
Respiration	40 ms	3.2	<250ms	Medium	Medium
Skin temperature	60 s	2.27	<250ms	Low	Medium
Glucose sensor	250 s	0.528	<250ms	Medium	Medium

Table 1: Traffic patterns and requirements of sensor nodes

3.2. Patient monitoring for Chronic Diseases

For a chronic disease patient, the formal procedure of routine visits is required to monitor the progress, development of complications or relapse of the disease. The questions like what, how and when to monitor are really crucial for disease treatment. In this context, various biosensors can be used for monitoring the patient's physiological conditions which brings relevant information on a regular basis. Appendix A and B shows patient monitoring use case scenario for WBAN.

3.3. Elderly Patient Monitoring

The fast growth in the elderly population will produce a considerable shortage of healthcare experts in the near future. WBAN delivers a highly cost effective solution to monitor the physiological parameters of the elderly persons by seamless integration of their daily routine activities. Moreover, the physician can monitor the health conditions of an elderly person remotely by the courtesy of WBANs.

4. Why 6lo is required for IEEE 802.15.6

Based on the characteristics defined in the overview section, the following sections elaborate on the main problems with IP for WBANs.

4.1. IPv6 Connectivity requirements

The requirement for IPv6 connectivity within WBANs is driven by the following:

- o The number of devices in WBANs makes network auto configuration and statelessness highly desirable. And for this, IPv6 has (default auto-configuration as a) ready solutions.
- o The large number of devices poses the need for a large address space, moreover a WBAN may consist of 256 nodes maximum and IPv6 is helpful to solve addressing issues.
- o Given the limited packet size of WBANs, the IPv6 address format allows subsuming of IEEE 802.15.6 addresses if so desired.
- o Simple interconnectivity to other IP networks including the Internet.
- o However, given the limited packet size, headers for IPv6 and layers above must be compressed whenever possible.

However, given the limited packet size, headers for IPv6 and layers above must be compressed whenever possible.

4.2. Limited Packet Size

Applications within WBANs are expected to originate small packets. Adding all layers for IP connectivity should still allow transmission in one frame, without incurring excessive fragmentation and reassembly. Furthermore, protocols must be designed or chosen so that the individual "control/protocol packets" fit within a single 802.15.6 frame. Along these lines, IPv6's requirement of sub-IP reassembly may pose challenges for low-end WBANs healthcare devices that do not have enough RAM or storage for a 1280-octet packet [RFC2460].

4.3. Topology requirements

The IEEE 802.15.6 working group has considered WBANs to operate in either a one-hop or two-hop star topology with the node in the centre of the star being placed on a location like the waist. Two feasible types of data transmission exist in the one-hop star topology: transmission from the device to the coordinator and transmission from the coordinator to the device. The communication methods that exist in the star topology are beacon mode and non-beacon mode. In a two-hop star WBAN, a relay-capable node may be used to exchange data frames between a node and the hub.

5. Scope/Purpose

This is a standard for short-range, wireless communication in the vicinity of, or inside, a human body (but not limited to humans). It uses existing industrial scientific medical (ISM) bands as well as frequency bands approved by national medical and/or regulatory authorities. Support for quality of service (QoS), extremely low power, and data rates from 10Kbps to 10 Mbps is required while simultaneously complying with strict non-interference guidelines where needed. The Table 1 shows a comparison of WBAN and other available technologies in terms of data rate and power consumption.

Standard	Provided data rate	Power requirement	Battery lifetime
802.11 ac (WiFi)	700 Mbps	100 mW - 1000 mW	Hours - days
Bluetooth	1Mbps - 10 Mbps	4 mW - 100 mW	Days - weeks
Wibree	600 Kbps maximum	2 mW - 10 mW	Weeks - months
ZigBee	250 Kbps	3 mW - 10 mW	Weeks - months
802.15.4	250 Kbps maximum	3 mW - 10 mW	Weeks - months
802.15.6	1Kbps - 10 Mbps	0.1 mW - 2 mW	Months - years

Table 2: Comparison of WBAN

The purpose of this document is to provide an international standard for a short-range (i.e., about human body range), low power, and highly reliable wireless communication for use in close proximity to, or inside, a human body. Data rates, typically up to 10Mbps, can be offered to satisfy an evolutionary set of entertainment and healthcare services. Current personal area networks (PANs) do not meet the medical (proximity to human tissue) and relevant communication regulations for some application environments. They also do not support the combination of reliability, QoS, low power, data rate, and non-interference required to broadly address the breadth of body area network (BAN) applications.

6. Layer 2 Overview

All nodes and hubs (coordinator in 802.15.4) are to be organized into logical sets, referred to as body area networks (BANs) in this specification, and coordinated by their respective hubs for medium access and power management as illustrated in Table 1. There is to be one and only one hub in a BAN, whereas the number of nodes in a BAN is to range from zero to mMaxBANSize. In a one-hop star BAN [SURVEY-WBAN][RFC7326], frame exchanges are to occur directly between nodes and the hub of the BAN. In a two-hop extended star BAN, the hub and a node are to exchange frames optionally via a relay-capable node. Some of the characteristics of WBANs are as follows:

6.1. Frame format

Figure 1 shows the general MAC frame format consisting of a 56-bit header, variable length frame body, and 18-bit FrameCheck Sequence (FCS). The maximum length of the frame body is 255 octets. The MAC header further consists of 32-bit frame control, 8-bit recipient Identification (ID), 8-bit sender ID, and 8-bit WBAN ID fields. The frame control field carries control information including the type of frame, that is, beacon, acknowledgement, or other control frames. The recipient and sender ID fields contain the address information of the recipient and the sender of the data frame, respectively. The WBAN ID contains information on the WBAN in which the transmission is active. The first 8-bit field in the MAC frame body carries message freshness information required for nonce construction and replay detection. The frame payload field carries data frames, and the last 32-bit Message Integrity Code (MIC) carries information about the authenticity and integrity of the frame.

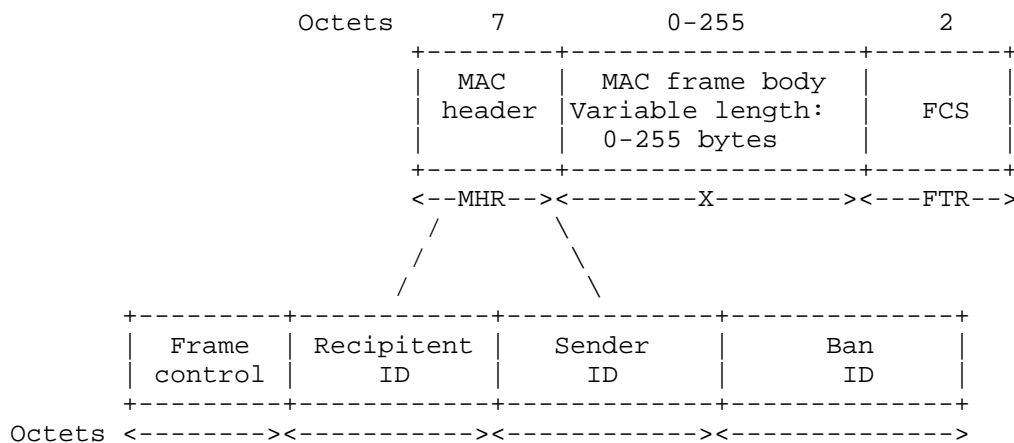


Figure 1: The general MAC frame format of IEEE 802.15.6

6.2. Frequency bands

The USA Federal Communications Commission (FCC) and communication authorities of other countries have allocated the MICS band at 402-405 MHz with 300 KHz channels to enable wireless communication with implanted medical devices [[[REFERENCE TO BE ADDED]]]. This leads to better penetration through the human tissue compared to higher frequencies, high level of mobility, comfort and better patient care in implant to implant (S1), implant to body surface (S2) and implant to external (S3) scenarios. Additionally, the 402-405 MHz frequencies offers conducive propagation characteristics for the transmission of radio signals in the human body and do not cause

severe interference for other radio operations in the same band. In fact, the MICs band is an unlicensed, ultra-low power, mobile radio service for transmitting data to support therapeutic or diagnostic operation related to implant medical devices and is internationally available. It is specifically chosen to provide low-power, small size, fast data transfer as well as a long communication range [SURVEY-WBAN][MAC-WBAN]. The frequency range of the MICS band allows high-level integration with the radio frequency IC (RFIC) technology, which leads to miniaturization and low power consumption. The PHY layer of IEEE 802.15.6 is responsible for the following tasks: activation and deactivation of the radio transceiver, Clear channel assessment (CCA) within the current channel and data transmission and reception. The choice of the physical layer depends on the target application: medical/non-medical, in, on and off-body. The PHY layer provides a procedure for transforming a physical layer service data unit (PSDU) into a physical layer protocol data unit (PPDU). IEEE 802.15.6 has specified three different physical layers: Human Body Communication (HBC), Narrow Band (NB) and Ultra-Wide Band (UWB). Various frequency bands are supported and shown in Table 2.

Communication	Frequency	Bandwidth
HBC	16 MHz	4 MHz
HBC	27 MHz	4 MHz
NB	402-405 MHz	300 KHz
NB	420-450 MHz	300 KHz
NB	863-870 MHz	400 KHz
NB	902-928 MHz	500 KHz
NB	956-956 MHz	400 KHz
NB	2360-2400 MHz	1 MHz
NB	2400-2438.5 MHz	1 MHz
UWB	13.2-4.7 GHz	499 MHz
UWB	6.2-10.3 GHz	499 MHz

Table 3: Frequency bands and Channel bandwidth of IEEE 802.15.6

6.3. Channel modes of IEEE 802.15.6

o Beacon Mode with Beacon Period Superframe Boundaries:

Beacons are sent at beacon periods by the coordinator and the superframe structure is managed by the coordinator by using beacon frames. The Physical Protocol Data Unit (PPDU) frame of Narrowband (NB) consists of a PHY Service Data Unit (PSDU) and Physical Layer Convergence Procedure (PLCP). PLCP preamble supports the receiver for synchronization process and considers as first module being send at given symbol rate. PLCP header sends decoding information for the receiver and it is transmitted after PLCP preamble. PSDU is last module of PPDU and consists of MAC header, Frame Check Sequence (FCS) and MAC frame body. PSDU is transmitted after PLCP with help of available frequency band with specific data rates. Different modulations schemes can be used with NB, namely, Differential Binary Phase-shift Keying (DBPSK), Differential Quadrature Phase-shift Keying (DQPSK) and Differential 8-Phase-shift Keying (D8PSK). NB uses seven frequency bands and operates under different data rates and modulation schemes. Medical Implant Communication Service (MICS) is the first licensed band of NB and used for implant communication with range of 402-405 MHz in most countries. Lower frequencies possess less attenuation and shadowing effect from body. Wireless Telemetry Medical Services (WMTS) is another license band and used for telemetry services. Although, Industrial, Scientific and Medical (ISM) band is free worldwide but it generates high probability of interference for IEEE 802.15.4 and IEEE 802.15.6 devices and considered as 7th license-free band. The 6th band (2360-2400 MHz) is used for medical devices instead of ISM band and offers less interference.

The superframe structure consists of several phases: exclusive access phase 1 (EAP 1), random access phase 1 (RAP1), type I/II phase, an EAP 2, RAP 2 and contention access phase (CAP). CSMA/CA or slotted Aloha is used by EAPs, RAPs and CAPs. For emergency services and high priority data, the EAP 1 and EAP 2 are used, whereas, CAP, RAP 1 and RAP 2 are used for regular data traffic. Type I/II are used for bi-link allocation intervals, up-link and down-link allocation intervals and delay bi-link intervals. For resource allocation, the type I/II polling is used.

A node's backoff counter value is set to a random integer number in the range $[1, CW \text{ (Contention Window)}]$, where CW (default value is CW_{min}) belongs to CW_{min} and CW_{max} which is dependent on user priority. When the algorithm starts, node begins counter decrement by one for every idle CSMA/CA slot duration (slot duration is equal to Pcsma/CA slot length). A node considers a CSMA/CA slot idle if the channel has been idle between start of slot and $pCCATime$. When

the backoff counter reaches zero, the node transmits the data frame. In case the channel is busy because of some other frame transmission, then node locks its backoff counter until the channel gets idle. The value of CW get double in case of even number of failures until it reaches CWmax [CHALLENGES-WBAN] [RFC7548].

o Beacon Mode with Superframe Boundaries:

For this mode, the coordinator provides an unscheduled polled allocation and each node establishes its own schedule. Different access mechanisms are used in superframe phases: schedule access (connection oriented and contention-free access), improvised and unscheduled access (connectionless and contention free access) and random access (CSMA/CA or slotted Aloha based).

o Beacon Mode without Superframe Boundaries:

In this channel access mode, beacons are not transmitted and channel is assigned by using polling mechanism.

7. --IETF to standardize--

This draft intend to standardize IEEE 802.15.6 for WBANs, specifically for implantable and wearable sensors. By standardizing means that integration of frame format need to be done i.e., how the IEEE 802.15.6 frame format will communicate with IPv6? How 6LoWPAN can accommodate this different frame format? The purpose of the mentioned use cases is to highlight the importance of the standard.

The 6LoWPAN is used to provide integration between IEEE 802.15.4 and IPv6. The details are mentioned in [RFC7548]. The 6LoWPAN concept originated with the purpose of connectivity of internet protocol with low-power smaller devices so they could claim to be part of Internet of Things (IoT) Networks.

The 6LoWPAN group has defined encapsulation and header compression mechanisms that allow ipv6 packets to be sent and received over IEEE 802.15.4 based networks, similarly the draft intent to define these mechanisms for IEEE 802.15.6. The 6LoWPAN can not be used with IEEE 802.15.6 due to frame size differences of IEEE 802.15.4 and IEEE 802.15.6.

8. IANA Considerations

[TBD]

9. Security Considerations

IPv6 over WBAN's applications often require confidentiality and integrity protection. This can be provided at the application, transport, network, and/or at the link.

10. Acknowledgements

[TBD]

11. References

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC7548] Ersue, M., Ed., Romascanu, D., Schoenwaelder, J., and A. Sehgal, "Management of Networks with Constrained Devices: Use Cases", RFC 7548, DOI 10.17487/RFC7548, May 2015, <<http://www.rfc-editor.org/info/rfc7548>>.
- [RFC7326] Parello, J., Claise, B., Schoening, B., and J. Quittek, "Energy Management Framework", RFC 7326, DOI 10.17487/RFC7326, September 2014, <<http://www.rfc-editor.org/info/rfc7326>>.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <<http://www.rfc-editor.org/info/rfc2460>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<http://www.rfc-editor.org/info/rfc4944>>.

11.2. Informative References

- [I-D.ietf-6tisch-6top-sf0] Dujovne, D., Grieco, L., Palattella, M., and N. Accettura, "6TiSCH 6top Scheduling Function Zero (SF0)", draft-ietf-6tisch-6top-sf0-02 (work in progress), October 2016.

[I-D.satish-6tisch-6top-sf1]

Anamalamudi, S., Zhang, M., Sangi, A., Perkins, C., and S. Anand, "Scheduling Function One (SF1) for hop-by-hop Scheduling in 6tisch Networks", draft-satish-6tisch-6top-sf1-02 (work in progress), August 2016.

[IEEE802.15.6]

"IEEE Standard, 802.15.6-2012 - IEEE Standard for Local and metropolitan area networks - Part 15.6: Wireless Body Area Networks", 2012,
<<https://standards.ieee.org/findstds/standard/802.15.6-2012.html>>.

[SURVEY-WBAN]

Diffie, W., Samaneh Movassaghi, Mehran Abolhasan, Justin Lipman, David Smith, and Abbas Jamalipour, "Wireless body area networks: A survey", Communications Surveys and Tutorials, IEEE , vol. 16, no. 3, pp. 1658-1686, 2014.

[MAC-WBAN]

Minglei Shu, Dongfeng Yuan, Chongqing Zhang, Yinglong Wang, and Changfang Chen, "A MAC Protocol for Medical Monitoring Applications of Wireless Body Area Networks.", Sensors , vol. 15, no. 6, 2015.

[CHALLENGES-WBAN]

Riccardo Cavallari, Flavia Martelli, Ramona Rosini, Chiara Buratti, and Roberto Verdone, "A Survey on Wireless Body Area Networks: Technologies and Design Challenges.", IEEE Communications Surveys and Tutorials , vol. 16, no. 3, pp. 1635-1657, 2014.

Appendix A. Patient monitoring use case - Spoke Hub

Refer following diagram:




Figure 2: Patient monitoring use case - Spoke Hub

Appendix B. Patient monitoring use case - Connected

Refer following diagram:




Figure 3: Patient monitoring use case - Connected

Authors' Addresses

Muhammad Sajjad Akbar
Bournemouth University
Fern Barrow, Dorset
Poole BH12 5BB
United Kingdom

Email: makbar@bournemouth.ac.uk

Naveed Bin Rais
Ajman University
University Street, Al jerf 1
Ajman 346
United Arab Emirates

Email: naveedbinrais@gmail.com

Abdur Rashid Sangi
Individual Contributor

Email: sangi_bahrian@yahoo.com

Mingui Zhang
Huawei Technologies
No. 156 Beiqing Rd. Haidian District
Beijing 100095
China

Email: zhangmingui@huawei.com

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: October 8, 2017

P. Thubert, Ed.
Cisco Systems
J. Hui
Nest Labs
April 6, 2017

LLN Fragment Forwarding and Recovery
draft-thubert-6lo-forwarding-fragments-05

Abstract

Considering that an LLN link-layer frame can have a payload below 100 bytes, an IPv6 packet might be fragmented more than 10 fragments at the 6LoWPAN layer. In a 6LoWPAN mesh-under mesh network, the fragments can be forwarded individually across the mesh, whereas a route-over mesh network, a fragmented 6LoWPAN packet must be reassembled at every hop, which causes latency and congestion. This draft introduces a simple protocol to forward individual fragments across a route-over mesh network, and, regardless of the type of mesh, recover the loss of individual fragments across the mesh and protect the network against bloat with a minimal flow control.

Status of This Memo

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

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

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

This Internet-Draft will expire on October 8, 2017.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

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

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

Table of Contents

1. Introduction	2
2. Updating RFC 4944	3
3. Terminology and Referenced Work	4
4. New Dispatch types and headers	5
4.1. Recoverable Fragment Dispatch type and Header	5
4.2. RFRAG Acknowledgment Dispatch type and Header	6
5. Fragments Recovery	8
6. Forwarding Fragments	9
6.1. Upon the first fragment	10
6.2. Upon the next fragments	11
6.3. Upon the RFRAG Acknowledgments	11
7. Security Considerations	12
8. IANA Considerations	12
9. Acknowledgments	12
10. References	12
10.1. Normative References	12
10.2. Informative References	13
Appendix A. Rationale	14
Appendix B. Requirements	16
Appendix C. Considerations On Flow Control	17
Authors' Addresses	18

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 10K bytes 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 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. Those undesirable effects can be alleviated by a hop-by-hop fragment forwarding technique such as the one proposed in this specification, and arguably this could be achieved 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 uncontrolled state in the network.

Another issue against RFC 4944 [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. Updating RFC 4944

This specification deprecates the fragmentation mechanism that is specified in RFC 4944 [RFC4944] and replaces it with a model where fragments can be forwarded end-to-end across a 6LoWPAN mesh network of any type, and where fragments that are lost on the way can be recovered individually. New dispatch types are defined in Section 4.

3. Terminology and 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 RFC 4963 [RFC4963] and follow the references for more information.

That experience led to the definition of "Path MTU discovery" [RFC1191] (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" [RFC5405].

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" [RFC8087]: 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". That technique is leveraged in this specification to forward fragments that actually do not have a network layer header, since the fragmentation occurs below IP.

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.

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].

4. New Dispatch types and headers

This specification aims at enabling to provide an MTU that is equivalent to 2048 bytes to the upper layer, which can be the 6LoWPAN Header Compression 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 and allows to forward fragments across a 6LoWPAN route-over mesh, but is not needed in the mesh-under case. The `datagram_tag` is used as the label and is locally unique to the node that is the MAC-layer source of the fragment. There is thus no need for a global registry of `datagram_tags` and a node may build the `datagram_tag` in its own locally-significant way, as long as the resulting tag stays unique to the particular datagram for the lifetime of that datagram.

This specification extends RFC 4944 [RFC4944] with 4 new Dispatch types, for Recoverable Fragments (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-ECN).

(to be confirmed by IANA) The new 6LoWPAN Dispatch types use the Value Bit Pattern of 11 1010xx, 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-ECN - RFRAG Ack with ECN Echo

Figure 1: Additional Dispatch Value Bit Patterns

4.1. Recoverable Fragment Dispatch type and Header

In this specification, the size and offset of the fragments are expressed on the compressed packet per 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`.

An RFRAG Acknowledgment Bitmap, whereby but at offset x indicates that fragment x was received.

5. 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.

The 6LoWPAN endpoint that fragments the packets at 6LoWPAN level (the sender) also controls the RFRAG Acknowledgments by setting the Ack Requested flag in the RFRAG packets. It may set the Ack Requested flag on any fragment so as to implement its own policy or perform congestion control by limiting the number of fragments in the air, IOW fragments that have been sent but for which reception or loss was not positively confirmed by the other 6LoWPAN endpoint. When the sender of the fragment knows that an underlying link-layer mechanism protects the Fragments already 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, though it may slightly defer it to let additional packets in.

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 interacts with the sender using an Acknowledgment message with a bitmap that indicates which fragments were actually received. The bitmap is a 32bit bitstring (a DWORD), which accommodates up to 32 fragments and is sufficient to transport 2028 bytes over an IEEE Std. 802.15.4 MAC payload. For all n in [0..31], bit n is set to 1 in the bitmap to indicate that fragment with

sequence *n* was received, otherwise the bit is set to 0. All 0s is a NULL bitmap that indicates that the fragmentation process was canceled by the receiver for that datagram.

The receiver MAY issue unsolicited acknowledgments. An unsolicited acknowledgment enables the sender endpoint to resume sending if it had reached its maximum number of outstanding fragments or indicate that the receiver has cancelled the process of an individual datagram. Note that acknowledgments might consume precious resources so the use of unsolicited acknowledgments should be configurable and not enabled by default.

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

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

When the sender decides that a packet should be dropped and the fragmentation process canceled, it sends a pseudo fragment with the `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. Upon an acknowledgment with a NULL bitmap, the sender MUST drop the datagram.

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

6.1. Upon the first fragment

In Route-Over mode, the MAC address changes at each hop. The label that is formed and placed in the datagram_tag is associated to the source MAC and only valid (and unique) for that source MAC. Say the first fragment has:

Source IPv6 address = IP_A (maybe hops away)

Destination IPv6 address = IP_B (maybe hops away)

Source MAC = MAC_prv (prv as previous)

Datagram_tag= DT_prv

The intermediate router that forwards individual fragments does the following:

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

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

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

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

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

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

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

swaps the MAC info to from self to MAC_nxt

Swaps the datagram_tag to DT_nxt

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

6.2. Upon the next fragments

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

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

swaps the MAC info to from self to MAC_nxt;

Swaps the datagram_tag to DT_nxt

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

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

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

Swaps the datagram_tag set to DT_prv

Bitmap of all 0es to indicate the error

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

6.3. Upon the RFRAG Acknowledgments

Upon an RFRAG Acknowledgment, the router expects to have already Label swap structure indexed by (MAC_nxt, DT_nxt), which are respectively the source MAC address of the received frame and the received datagram_tag. DT_nxt should have been computed by this router and this router should have assigned it to this particular datagram. The router:

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

swaps the MAC info to from self to MAC_prv;

Swaps the datagram_tag to DT_prv

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

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

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

7. 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].

8. IANA Considerations

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

9. Acknowledgments

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

10. References

10.1. Normative References

- [IEEE.802.15.4]
IEEE, "IEEE Standard for Low-Rate Wireless Networks",
IEEE Standard 802.15.4,
<<http://ieeexplore.ieee.org/document/7460875/>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", BCP 14, RFC 2119,
DOI 10.17487/RFC2119, March 1997,
<<http://www.rfc-editor.org/info/rfc2119>>.

- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<http://www.rfc-editor.org/info/rfc4944>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<http://www.rfc-editor.org/info/rfc6282>>.
- [RFC6298] Paxson, V., Allman, M., Chu, J., and M. Sargent, "Computing TCP's Retransmission Timer", RFC 6298, DOI 10.17487/RFC6298, June 2011, <<http://www.rfc-editor.org/info/rfc6298>>.

10.2. 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-11 (work in progress), January 2017.
- [RFC1191] Mogul, J. and S. Deering, "Path MTU discovery", RFC 1191, DOI 10.17487/RFC1191, November 1990, <<http://www.rfc-editor.org/info/rfc1191>>.
- [RFC2309] Braden, B., Clark, D., Crowcroft, J., Davie, B., Deering, S., Estrin, D., Floyd, S., Jacobson, V., Minshall, G., Partridge, C., Peterson, L., Ramakrishnan, K., Shenker, S., Wroclawski, J., and L. Zhang, "Recommendations on Queue Management and Congestion Avoidance in the Internet", RFC 2309, DOI 10.17487/RFC2309, April 1998, <<http://www.rfc-editor.org/info/rfc2309>>.
- [RFC2914] Floyd, S., "Congestion Control Principles", BCP 41, RFC 2914, DOI 10.17487/RFC2914, September 2000, <<http://www.rfc-editor.org/info/rfc2914>>.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, DOI 10.17487/RFC3031, January 2001, <<http://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, <<http://www.rfc-editor.org/info/rfc3168>>.

- [RFC4919] Kushalnagar, N., Montenegro, G., and C. Schumacher, "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals", RFC 4919, DOI 10.17487/RFC4919, August 2007, <<http://www.rfc-editor.org/info/rfc4919>>.
- [RFC4963] Heffner, J., Mathis, M., and B. Chandler, "IPv4 Reassembly Errors at High Data Rates", RFC 4963, DOI 10.17487/RFC4963, July 2007, <<http://www.rfc-editor.org/info/rfc4963>>.
- [RFC5405] Eggert, L. and G. Fairhurst, "Unicast UDP Usage Guidelines for Application Designers", RFC 5405, DOI 10.17487/RFC5405, November 2008, <<http://www.rfc-editor.org/info/rfc5405>>.
- [RFC5681] Allman, M., Paxson, V., and E. Blanton, "TCP Congestion Control", RFC 5681, DOI 10.17487/RFC5681, September 2009, <<http://www.rfc-editor.org/info/rfc5681>>.
- [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, <<http://www.rfc-editor.org/info/rfc7554>>.
- [RFC8087] Fairhurst, G. and M. Welzl, "The Benefits of Using Explicit Congestion Notification (ECN)", RFC 8087, DOI 10.17487/RFC8087, March 2017, <<http://www.rfc-editor.org/info/rfc8087>>.

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:

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

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

From the LLN node:

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

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

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

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

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

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

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

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.

Support for out-of-order fragment delivery

Forwarding over a mesh network with rerouting and load balancing can introduce out-of-sequence packets.

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

Optional congestion control

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

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

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 4.2.

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 make full sense to fire a next fragment as soon as the previous fragment has had its chance to be forwarded at the next hop, retry (ARQ) operations included.

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

acknowledgment might be on the way back. It is also possible that either the fragment or the acknowledgment was lost on the way.

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

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

Section 5 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.

Authors' Addresses

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

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

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

Email: jonhui@nestlabs.com