6lo
Internet-Draft
Updates: 6775 (if approved)
Intended status: Standards Track
Expires: April 21, 2019

P. Thubert, Ed.
Cisco
B. Sarikaya
M. Sethi
Ericsson
R. Struik
Struik Security Consultancy

October 18, 2018

Address Protected Neighbor Discovery for Low-power and Lossy Networks
draft-ietf-6lo-ap-nd-08

Abstract

This document specifies an extension to 6LoWPAN Neighbor Discovery (ND) defined in RFC6775 and updated in [I-D.ietf-6lo-rfc6775-update]. 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.

Table of Contents

1. Introduction

Neighbor Discovery Optimizations for 6LoWPAN networks [RFC6775] (6LoWPAN ND) adapts the original IPv6 neighbor discovery (NDv6) protocols defined in [RFC4861] and [RFC4862] for constrained low-power and lossy network (LLN). In particular, 6LoWPAN ND introduces a unicast host address registration mechanism that reduces the use of multicast. 6LoWPAN ND defines a new Address Registration Option (ARO) that is carried in the unicast Neighbor Solicitation (NS) and Neighbor Advertisement (NA) messages exchanged between a 6LoWPAN Node (6LN) and a 6LoWPAN Router (6LR). It also defines the Duplicate Address Request (DAR) and Duplicate Address Confirmation (DAC) messages between the 6LR and the 6LoWPAN Border Router (6LBR). In LLN networks, the 6LBR is the central repository of all the registered addresses in its domain.

The registration mechanism in 6LoWPAN ND [RFC6775] prevents the use of an address if that address is already registered in the subnet (first come first serve). In order to validate address ownership, the registration mechanism enables the 6LR and 6LBR to validate the association between the registered address of a node, and its Registration Ownership Verifier (ROVR). ROVR is defined in [I-D.ietf-6lo-rfc6775-update] and it can be derived from the MAC address of the device (using the 64-bit Extended Unique Identifier EUI-64 address format specified by IEEE). However, the EUI-64 can be spoofed, and therefore, any node connected to the subnet and aware of a registered-address-to-ROVR mapping could effectively fake the ROVR. This would allow the an attacker to steal the address and redirect traffic for that address. [I-D.ietf-6lo-rfc6775-update] defines an Extended Address Registration Option (EARO) option that allows to transport alternate forms of ROVRs, and is a pre-requisite for this specification.

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

The protected address registration protocol proposed in this document enables Source Address Validation (SAVI) [RFC7039]. This ensures that only the actual owner uses a registered address in the IPv6 source address field. A 6LN can only use a 6LR for forwarding packets only if it has previously registered the address used in the source field of the IPv6 packet.

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

2. Terminology

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

2.1. References

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

- SEcure Neighbor Discovery (SEND) [RFC3971]
- Cryptographically Generated Addresses (CGA) [RFC3972]
- Neighbor Discovery for IP version 6 [RFC4861]
- IPv6 Stateless Address Autoconfiguration[RFC4862],
- Problem Statement and Requirements for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing [RFC6606]
- 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]
- Terms Used in Routing for Low-Power and Lossy Networks (LLNs) [RFC7102]
- Terminology for Constrained-Node Networks [RFC7228]
- Registration Extensions for 6LoWPAN Neighbor Discovery" [I-D.ietf-6lo-rfc6775-update]

2.2. 6LoWPAN sub-glossary

This document uses the following acronyms:

6BBR: 6LoWPAN Backbone Router (proxy for the registration)[I-D.ietf-6lo-backbone-router]

6LBR: 6LoWPAN Border Router

6LN: 6LoWPAN Node

6LR: 6LoWPAN Router (relay to the registration process)

CIPO: Crypto-ID Parameters Option

(E)ARO: (Extended) Address Registration Option

DAD: Duplicate Address Detection

LLN: Low-Power and Lossy Network (a typical IoT network)

NA: Neighbor Advertisement

ND: Neighbor Discovery

NDP: Neighbor Discovery Protocol

NDPSO: NDP Signature Option

NS: Neighbor Solicitation

ROVR: Registration Ownership Verifier (pronounced rover)

RA: Router Advertisement

RS: Router Solicitation

RSAO: RSA Signature Option

TID: Transaction ID (a sequence counter in the EARO)

3. Updating RFC 6775

This specification defines a cryptographic identifier (Crypto-ID) that can be used as a replacement to the MAC address in the ROVR field of the EARO option; the computation of the Crypto-ID is detailed in Section 4.1. A node in possession of the necessary cryptographic primitives SHOULD use Crypto-ID by default as ROVR in its registration. Whether a ROVR is a Crypto-ID is indicated by a new "C" flag in the NS(EARO) message.

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

4. New Fields and Options

In order to avoid the need for new ND option types, this specification reuses/ extends options defined in SEND [RFC3971] and 6LoWPAN ND [RFC6775] [I-D.ietf-6lo-rfc6775-update]. This applies in particular to the CGA option and the RSA Signature Option. This specification provides aliases for the specific variations of those options as used in this document. The presence of the EARO option in the NS/NA messages indicates that the options are to be processed as specified in this document, and not as defined in SEND [RFC3971].

4.1. New Crypto-ID

Each 6LN using this specification for address registration MUST support Elliptic Curve Crytpograhy (ECC) and a hash function. The choice of elliptic curves and hash function currently defined in this specification are listed in Section 8.2.

The Crypto-ID is computed by a 6LN as follows:

1. Depending on the Crypto-Type (see Section 8.2) used by the node, the hash function is applied to the JSON Web Key (JWK) [RFC7517] encoding of the public-key of the node.

2. The leftmost bits of the resulting hash, up to the size of the ROVR field, are used as the Crypto-ID.

4.2. Updated EARO

This specification updates the EARO option as follows:

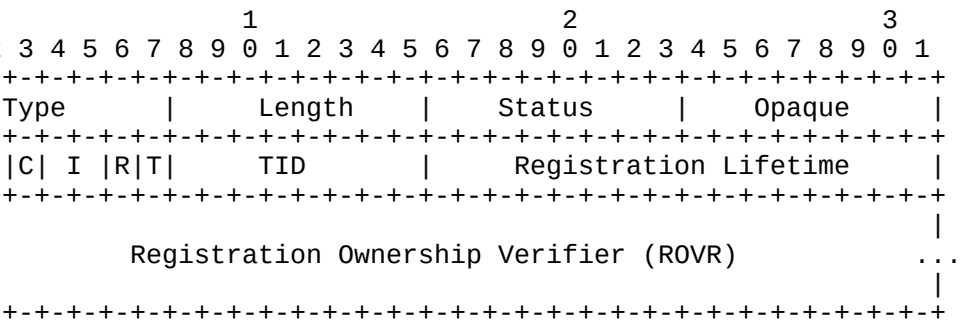

Figure 1: Enhanced Address Registration Option

Type: 33

Length: 8-bit unsigned integer. The length of the option (including the type and length fields) in units of 8 bytes.

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

Opaque: Defined in [I-D.ietf-6lo-rfc6775-update].

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

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

I: Defined in [I-D.ietf-6lo-rfc6775-update].

R: Defined in [I-D.ietf-6lo-rfc6775-update].

T and TID: Defined in [I-D.ietf-6lo-rfc6775-update].

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 6LoWPAN ND [I-D.ietf-6lo-rfc6775-update]. No other new Status values are defined.

4.3. Crypto-ID Parameters Option

This specification defines the Crypto-ID Parameters Option (CIPO), as a variation of the CGA Option that carries the parameters used to form a Crypto-ID. In order to provide cryptographic agility [RFC7696], AP-ND supports two possible elliptic curves, indicated by a Crypto-Type field. NIST P-256 [FIPS186-4] MUST be supported by all implementations. The Edwards-Curve Digital Signature Algorithm (EdDSA) curve Ed25519ph (pre-hashing) [RFC8032] MAY be supported as an alternate.

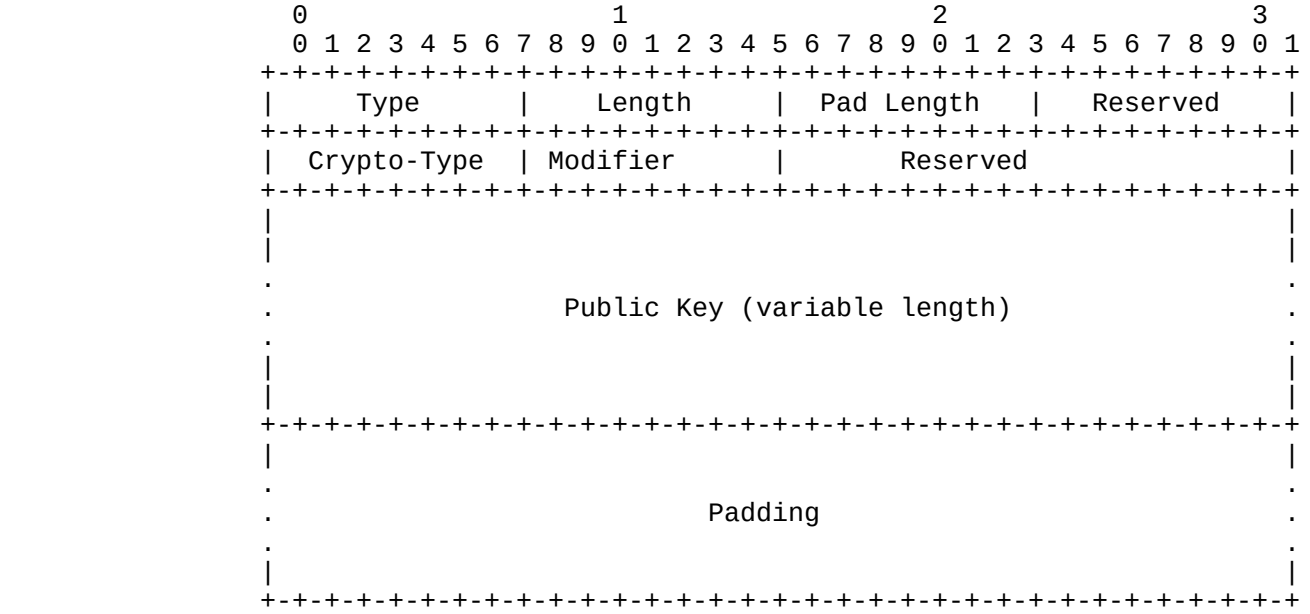

Figure 2: Crypto-ID Parameters Option

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

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

Modifier: 8-bit unsigned integer.

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

Crypto-Type: The type of cryptographic algorithm used in calculation Crypto-ID. A value of 0 indicates NIST

P-256, with SHA-256 as the hash algorithm. A value of 1 is assigned for Ed25519ph, with SHA-512 as the hash algorithm.

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

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

4.4. Nonce Option

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

4.5. NDP Signature Option

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

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

- The 128-bit CGA Message Type tag [RFC3972] for AP-ND is 0x8701 55c8 0cca dd32 6ab7 e415 f148 84d0. (The tag value has been generated by the editor of this specification on random.org).

- The signature is computed using the hash algorithm and the digital signature indicated in the Crypto-Type field of the CIPO option using the private-key corresponding the public-key passed in the CIPO.

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

5. Protocol Scope

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

The 6LBR maintains registration state for all devices in its attached LLN. Together with the first-hop router (the 6LR), the 6LBR assures uniqueness and grants ownership of an IPv6 address before it can be

used in the LLN. This is in contrast to a traditional network that relies on IPv6 address auto-configuration [RFC4862], where there is no guarantee of ownership from the network, and each IPv6 Neighbor Discovery packet must be individually secured [RFC3971].

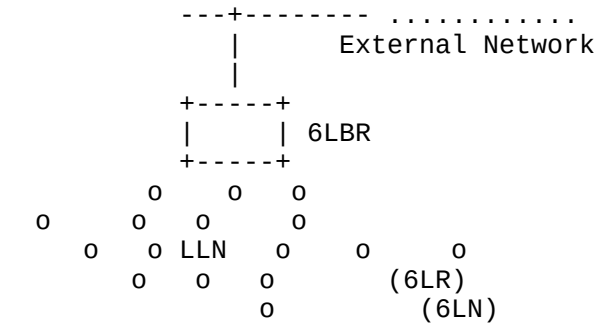

Figure 3: Basic Configuration

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

6. Protocol Flows

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

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

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

multiple addresses. The registration process allows the node to use the same Crypto-ID for all of its addresses.

6.1. First Exchange with a 6LR

A 6LN registers to a 6LR that is one hop away from it with the "C" flag set in the EARO, indicating that the ROVR field contains a Crypto-ID. The Target Address in the NS message indicates the IPv6 address that the 6LN is trying to register. The on-link (local) protocol interactions are shown in Figure 4. If the 6LR does not have a state with the 6LN that is consistent with the NS(EARO), then it replies with a challenge NA (EARO, status=Validation Requested) that contains a Nonce Option (shown as NonceLR in Figure 4). The Nonce option MUST contain a random Nonce value that was never used with this device.

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

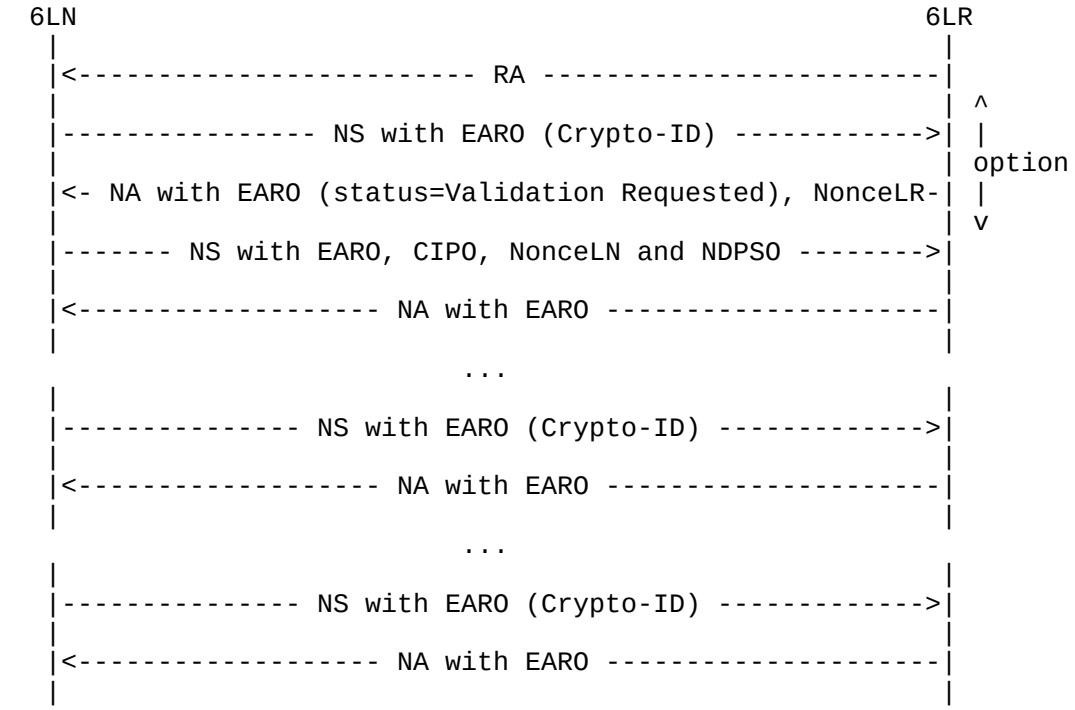

Figure 4: On-link Protocol Operation

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

- 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. The proof is not needed again in later registrations for that address. When a 6LR receives a NS(EARO) registration with a new Crypto-ID as a ROVR, it SHOULD challenge by responding with a NA(EARO) with a status of "Validation Requested".

- The challenge is triggered when the registration for a Source Link-Layer Address is not verifiable either at the 6LR or the 6LBR. In the latter case, the 6LBR returns a status of "Validation Requested" in the DAR/DAC exchange, which is echoed by the 6LR in the NA (EARO) back to the registering node. The challenge MUST NOT alter a valid registration in the 6LR or the 6LBR.

- Upon receiving a NA(EARO) with a status of "Validation Requested", the registering node SHOULD retry its registration with a Crypto-ID Parameters Option (CIPO) (Section 4.3) that contains all the

necessary material for building the Crypto-ID, the NonceLN that it generated, and the NDP signature (Section 4.5) option that proves its ownership of the Crypto-ID and intent of registering the Target Address.

- In order to validate the ownership, the 6LR performs the same steps as the 6LN and rebuilds the Crypto-ID based on the parameters in the CIPO. It also verifies the signature contained in the NDPSO option. If the Crypto-ID does not match with the public-key in the CIPO option, or if the signature in the NDPSO option cannot be verified, the validation fails.

- If the 6LR fails to validate the signed NS(EARO), it responds with a status of "Validation Failed". After receiving a NA(EARO) with a status of "Validation Failed", the registering node SHOULD try to register an alternate target address in the NS message.

6.2. NDPSO generation and verficiation

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 as follows:

- Concatenate the following in the order listed:

  1. 128-bit type tag (in network byte order)

  2. JWK-encoded public key

  3. the 16-byte Target Address (in network byte order) sent in the Neighbor Solicitation (NS) message. It is the address which the 6LN is registering with the 6LR and 6LBR.

  4. NonceLR received from the 6LR (in network byte order) in the Neighbor Advertisement (NA) message. The random nonce is at least 6 bytes long as defined in [RFC3971].

  5. NonceLN sent from the 6LN (in network byte order). The random nonce is at least 6 bytes long as defined in [RFC3971].

  6. The length of the ROVR field in the NS message cotainting the Crypto-ID that was sent.

  7. 1-byte (in network byte order) Crypto-Type value sent in the CIPO option.

- Depending on the Crypto-Type (see Section 8.2) chosen by the node (6LN), apply the hash function on this concatenation.

- Depending on the Crypto-Type (see Section 8.2) chosen by the node (6LN), sign the hash output with ECDSA (if curve P-256 is used) or sign the hash with EdDSA (if curve EdDSA25519ph).

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

- Concatenate the following in the order listed:

  1. 128-bit type tag (in network byte order)

  2. JWK-encoded public key received in the CIPO option

  3. the 16-byte Target Address (in network byte order) received in the Neighbor Solicitation (NS) message. It is the address which the 6LN is registering with the 6LR and 6LBR.

  4. NonceLR sent in the Neighbor Advertisement (NA) message. The random nonce is at least 6 bytes long as defined in [RFC3971].

  5. NonceLN received from the 6LN (in network byte order) in the NS message. The random nonce is at least 6 bytes long as defined in [RFC3971].

  6. The length of the ROVR field in the NS message containing the Crypto-ID that was received.

  7. 1-byte (in network byte order) Crypto-Type value received in the CIPO option.

- Depending on the Crypto-Type (see Section 8.2) indicated by the (6LN) in the CIPO, apply the hash function on this concatenation.

- Verify the signature with the public-key received and the locally computed values. If the verification succeeds, the 6LR and 6LBR add the state information about the Crypto-ID, public-key and Target Address being registered to their database.

6.3. Multihop Operation

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

A new device that joins the network auto-configures an address and performs an initial registration to a neighboring 6LR with an NS message that carries an Address Registration Option (EARO) [RFC6775]. The 6LR validates the address with an 6LBR using a DAR/DAC exchange, and the 6LR confirms (or denies) the address ownership with an NA message that also carries an Address Registration Option.

Figure 5 illustrates a registration flow all the way to a 6LowPAN Backbone Router (6BBR).

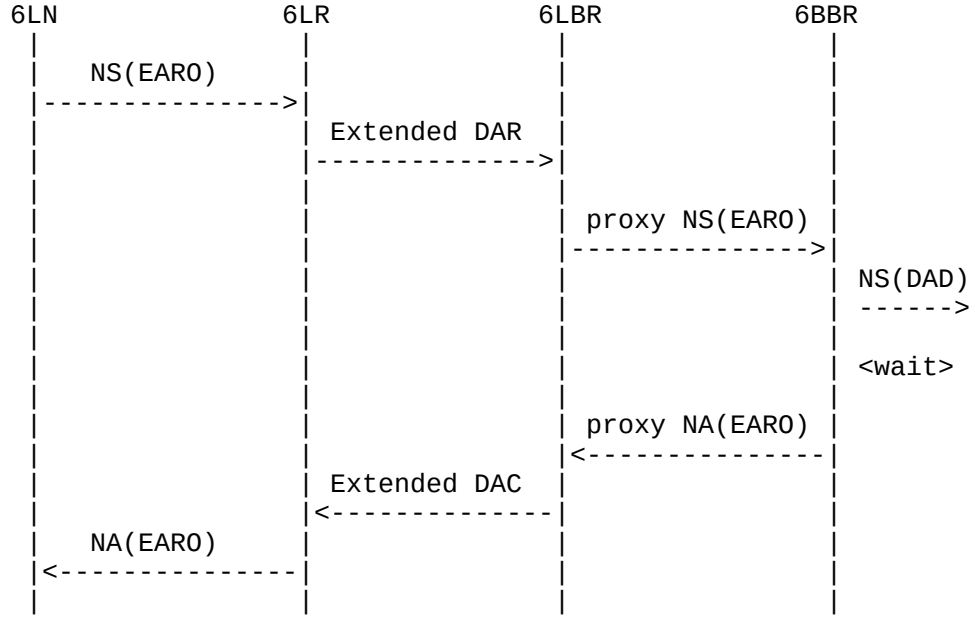

Figure 5: (Re-)Registration Flow

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

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

7. Security Considerations

7.1. Inheriting from RFC 3971

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

Neighbor Solicitation/Advertisement Spoofing

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

Duplicate Address Detection DoS Attack

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

Router Solicitation and Advertisement Attacks

This specification does not change the protection of RS and RA which can still be protected by SEND.

Replay Attacks

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

Neighbor Discovery DoS Attack

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

7.2. Related to 6LoWPAN ND

The threats discussed in 6LoWPAN ND [RFC6775] and its update [I-D.ietf-6lo-rfc6775-update] also apply here. Compared with SeND, this specification saves about 1Kbyte in every NS/NA message. Also, this specification separates the cryptographic identifier from the registered IPv6 address so that a node can have more than one IPv6 address protected by the same cryptographic identifier. SeND forces the IPv6 address to be cryptographic since it integrates the CGA as the IID in the IPv6 address. This specification frees the device to form its addresses in any fashion, thereby enabling not only 6LoWPAN compression which derives IPv6 addresses from Layer-2 addresses but also privacy addresses.

7.3. ROVR Collisions

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

8. IANA considerations

8.1. CGA Message Type

This document defines a new 128-bit value under the CGA Message Type [RFC3972] namespace, 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 0..255 and contains a Signature Algorithm and a Hash Function as shown in Table 1. The following Crypto-Type values are defined in this document:

Crypto-Type value	Signature Algorithm	Hash Function	Defining Specification
0	NIST P-256 [FIPS186-4]	SHA-256 [RFC6234]	RFC THIS
1	Ed25519ph [RFC8032]	SHA-512 [RFC6234]	RFC THIS

Table 1: Crypto-Types

As is evident from the table above, although the two curves provide similar security, they however rely on different hash functions. Supporting multiple hash functions on constrained devices is not ideal. [I-D.struik-lwig-curve-representations] provides information on how to represent Montgomery curves and (twisted) Edwards curves as curves in short-Weierstrass form and illustrates how this can be used to implement elliptic curve computations using existing implementations that already implement, e.g., ECDSA and ECDH using NIST [FIPS186-4] prime curves. New Crypto-Type values providing similar or better security (with less code) can be defined in future.

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

9. Acknowledgments

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

10. References

10.1. Normative References

[FIPS186-4]
FIPS 186-4, "Digital Signature Standard (DSS), Federal Information Processing Standards Publication 186-4", US Department of Commerce/National Institute of Standards and Technology Gaithersburg, MD, July 2013.

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

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

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

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

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <https://www.rfc-editor.org/info/rfc4861>.

[RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <https://www.rfc-editor.org/info/rfc4862>.

[RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <https://www.rfc-editor.org/info/rfc6775>.

[RFC7517] Jones, M., "JSON Web Key (JWK)", RFC 7517, DOI 10.17487/RFC7517, May 2015, <https://www.rfc-editor.org/info/rfc7517>.

10.2. Informative references

[I-D.ietf-6lo-backbone-router]
Thubert, P. and C. Perkins, "IPv6 Backbone Router", draft-ietf-6lo-backbone-router-07 (work in progress), September 2018.

[I-D.struik-lwig-curve-representations]
Struik, R., "Alternative Elliptic Curve Representations", draft-struik-lwig-curve-representations-02 (work in progress), July 2018.

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

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

[RFC6234] Eastlake 3rd, D. and T. Hansen, "US Secure Hash Algorithms (SHA and SHA-based HMAC and HKDF)", RFC 6234, DOI 10.17487/RFC6234, May 2011, <https://www.rfc-editor.org/info/rfc6234>.

[RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <https://www.rfc-editor.org/info/rfc6282>.

[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>.

[RFC7039] Wu, J., Bi, J., Bagnulo, M., Baker, F., and C. Vogt, Ed., "Source Address Validation Improvement (SAVI) Framework", RFC 7039, DOI 10.17487/RFC7039, October 2013, <https://www.rfc-editor.org/info/rfc7039>.

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

[RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <https://www.rfc-editor.org/info/rfc7217>.

[RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <https://www.rfc-editor.org/info/rfc7228>.

[RFC7696] Housley, R., "Guidelines for Cryptographic Algorithm Agility and Selecting Mandatory-to-Implement Algorithms", BCP 201, RFC 7696, DOI 10.17487/RFC7696, November 2015, <https://www.rfc-editor.org/info/rfc7696>.

[RFC8032] Josefsson, S. and I. Liusvaara, "Edwards-Curve Digital Signature Algorithm (EdDSA)", RFC 8032, DOI 10.17487/RFC8032, January 2017, <https://www.rfc-editor.org/info/rfc8032>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://www.rfc-editor.org/info/rfc8126>.

Appendix A. Requirements Addressed in this Document

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

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

Authors' Addresses

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

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

Behcet Sarikaya
Plano, TX
USA

Email: sarikaya@ieee.org

Mohit Sethi
Ericsson
Jorvas 02420
Finland

Email: mohit@piuha.net

Rene Struik
Struik Security Consultancy

Email: rstruik.ext@gmail.com

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

P. Thubert, Ed.
cisco
C. Perkins
Futurewei

October 22, 2018

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

Abstract

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

Status of This Memo

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

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

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

This Internet-Draft will expire on April 25, 2019.

Table of Contents

1. Introduction

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

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

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

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

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

These problems suggest restricting IPv6 ND broadcasts over wireless access links, which can be done by dividing up the subnet. Another way is to take over (proxy) the Layer-3 protocols that rely on broadcast operation at the boundary of the wired and wireless domains, emulating the Layer-2 association at layer-3. For instance, IEEE 802.11 [IEEEstd80211] specifies ARP and ND proxy [RFC4389] services at the Access Points (APs).

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

WPAN devices (i.e., those implementing IEEE STD. 802.15.4 [IEEEstd802154]) can make use of Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs) [RFC6775] which treats the wireless medium as different from Ethernet. RFC 6775 is updated as [I-D.ietf-6lo-rfc6775-update]; the update includes changes that are required by this document.

2. Applicability and Requirements Served

This specification updates and generalizes 6LoWPAN ND to a broader range of Low power and Lossy Networks (LLNs) with support for Duplicate Address Detection (DAD) and address lookup that does not require broadcasts over the LLNs. The term LLN is used loosely in this specification to cover multiple types of WLANs and WPANs, including Low-Power Wi-Fi, BLUETOOTH(R) Low Energy, IEEE STD. 802.11AH and IEEE STD. 802.15.4 wireless meshes, so as to address the requirements listed in Appendix B.3 of [I-D.ietf-6lo-rfc6775-update] "Requirements Related to the Variety of Low-Power Link types".

For the TimeSlotted Channel Hopping (TSCH) mode of [IEEEstd802154], the 6TiSCH architecture [I-D.ietf-6tisch-architecture] describes how a 6LoWPAN ND host could connect to the Internet via a RPL mesh Network, but doing so requires extensions to the 6LOWPAN ND protocol to support mobility and reachability in a secure and manageable environment. The extensions detailed in this document also work for the 6TiSCH architecture, serving the requirements listed in Appendix B.2 of [I-D.ietf-6lo-rfc6775-update] "Requirements Related to Routing Protocols".

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

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

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

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

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

3. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] .

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

- "Neighbor Discovery for IP version 6" [RFC4861],
- "IPv6 Stateless Address Autoconfiguration" [RFC4862],
- "Multi-Link Subnet Issues" [RFC4903],
- "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],
- ,"Mobility Support in IPv6" [RFC6275],
- "Neighbor Discovery Proxies (ND Proxy)" [RFC4389]
- "Optimistic Duplicate Address Detection" [RFC4429], and

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

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

Sleeping Proxy

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

Unicasting Proxy

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

Routing proxy

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

Bridging proxy

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

Primary BBR

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

Secondary BBR

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

4. Overview

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

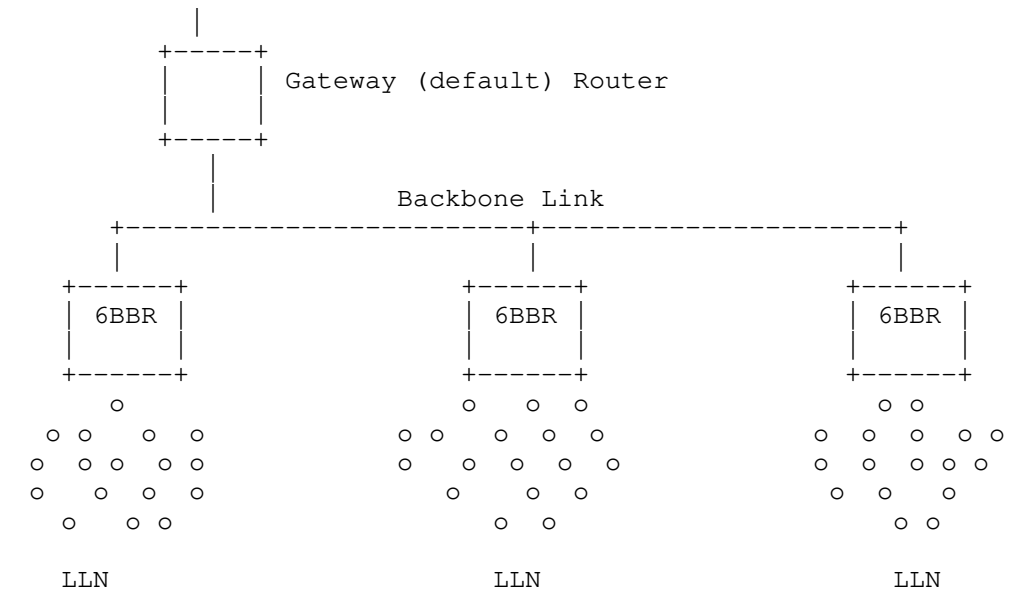

Figure 1: Backbone Link and Backbone Routers

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

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

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

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

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

5. Backbone Router Routing Operations

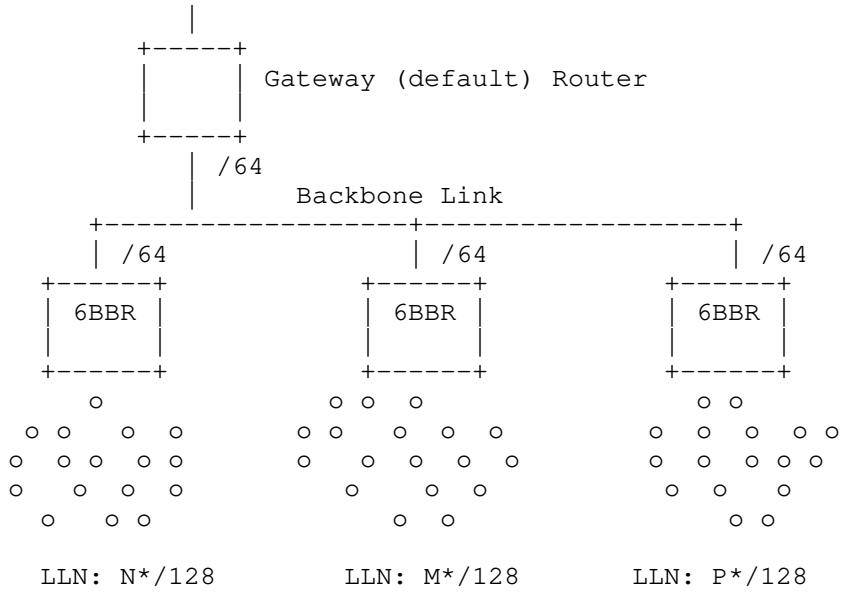

Figure 2: Example Routing Configuration for 3 LLNs in the ML Subnet

5.1. Over the Backbone Link

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

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

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

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

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

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

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

5.2. Proxy Operations Over the LLN Interface

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

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

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

5.2.1. Routing Proxy Operations

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

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

For each Registered Address, the 6BBR SHOULD maintain a list of the peers on the backbone which have associated its MAC address with the Registered Address. If that Registered Address moves to a different 6BBR, the first 6BBR SHOULD unicast a gratuitous NA(Override) to each such peer, to supply the MAC address of the new 6BBR in the TLLA option for the Address.

5.2.2. Bridging Proxy Operations

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

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

6. Backbone Router Proxy Operations

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

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

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

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

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

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

6.1. Primary and Secondary BBRs

A 6BBR MAY be primary or secondary. The primary is the backbone router that has the highest EUI-64 address of all the 6BBRs that share a registration for a same Registered Address, with the same ROVR and same Transaction ID, the EUI-64 address being considered as an unsigned 64bit integer. A given 6BBR can be primary for a given address and secondary for another address, regardless of whether or not the addresses belong to the same 6LN. The primary Backbone Router is in charge of protecting the address for DAD over the Backbone. Any of the Primary and Secondary 6BBR may claim the address over the backbone, since they are all capable to route from the backbone to the 6LN; the address appears on the backbone as an anycast address.

6.2. Binding Table

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

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

Only one 6LN may register a given Address at a particular 6BBR. However, that Registered Address may be registered to Multiple 6BBRs for higher availability.

Over the LLN, Binding Table management is as follows:

De-registrations (newer TID, same ROVR, null Lifetime) are accepted and acknowledged with a status of 4 (TBD); the entry is deleted;

Newer registrations (newer TID, same ROVR, non-null Lifetime) are acknowledged with a status of 0 (success); the binding is updated with the new TID, the Registration Lifetime and the Registering Node; in TENTATIVE state the acknowledgement is held and may be

overwritten; in other states the Registration-Lifetime timer is restarted and the entry is placed in REACHABLE state.

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

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

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

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

6.3. Registration and Binding Table Entry Creation

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

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

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

2. The entry is also removed if an NA with an ARO option with a status of 3 (moved), or a NS with an ARO option that indicates a newer registration for the same Registered Node, is received over

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

3. When a registration is updated but not deleted, e.g. from a newer registration, the DAD process on the backbone continues and the running timers are not restarted;

4. Other NS (including DAD with no EARO) and NA from the backbone are not acknowledged in TENTATIVE state. To cover legacy 6LNs that do not support ODAD, the list of their origins MAY be stored and then, if the TENTATIVE_DURATION timer elapses, the 6BBR MAY send each such legacy 6LN a unicast NA.

5. When the TENTATIVE_DURATION timer elapses, a status 0 (success) is returned in a NA(EARO) back to the Registering Node(s), and the entry goes to REACHABLE state for the Registration Lifetime. The 6BBR MUST send a multicast NA(EARO) to the SNMA associated to the Registered Address over the backbone with the Override bit set so as to take over the binding from other 6BBRs.

6.4. Defending Addresses

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

- If the 6BBR is primary, or does not support the function of primary, it MUST defend that address over the backbone upon receiving NS, either if the NS does not carry an EARO, or if an EARO is present that indicates a different Registering Node (different ROVR). The 6BBR sends a NA message with the Override bit set and the NA carries an EARO if and only if the NS-DAD did so. When present, the EARO in the NA(Override) that is sent in response to the NS(EARO) carries a status of 1 (duplicate), and the ROVR and TID fields in the EARO are obfuscated with null or random values to avoid network scanning and impersonation attacks.

- If the 6BBR receives an NS(EARO) for a newer registration, the 6BBR updates the entry and the routing state to forward packets to the new 6BBR, but keeps the entry REACHABLE. Afterwards, the 6BBR MAY use REDIRECT messages to reroute traffic for the Registered Address to the new 6BBR.

- If the 6BBR receives an NA(EARO) for a newer registration, the 6BBR removes its entry and sends a NA(EARO) with a status of 3 (MOVED) to the Registering Node, if the Registering Node is different from the Registered Node. The 6BBR cleans up existing Neighbor Cache entries in peer nodes as discussed in Section 5.1, by unicasting to each such peer, or one broadcast NA(Override).

- If the 6BBR receives a NS(LOOKUP) for a Registered Address, it answers immediately with an NA on behalf of the Registered Node, without polling it. There is no need of an EARO in that exchange.

- When the Registration-Lifetime timer elapses, the entry goes to STALE state for a duration of STABLE_STALE_DURATION in LLNs that keep stable addresses such as LWPANs, and UNSTABLE_STALE_DURATION in LLNs where addresses are renewed rapidly, e.g. for privacy reasons.

The STALE state enables tracking of the backbone peers that have a Neighbor Cache entry pointing to this 6BBR in case the Registered Address shows up later. If the Registered Address is claimed by another 6LN on the backbone, with an NS-DAD or an NA, the 6BBR does not defend the address. In STALE state:

- If STALE_DURATION elapses, the 6BBR removes the entry.

- Upon receiving an NA(Override) the 6BBR removes its entry and sends a NA(EARO) with a status of 4 (removed) to the Registering Node.

- If the 6BBR receives a NS(LOOKUP) for a Registered Address, the 6BBR MUST send an NS(NUD) following rules in [RFC7048] to the Registering Node targeting the Registered Address prior to answering. If the NUD succeeds, the operation in REACHABLE state applies. If the NUD fails, the 6BBR refrains from answering the lookup. The NUD SHOULD be used by the Registering Node to indicate liveness of the Registered Node, if they are different nodes.

7. Security Considerations

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

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

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

8. Protocol Constants

This Specification uses the following constants:

TENTATIVE_DURATION: 800 milliseconds

STABLE_STALE_DURATION: 24 hours

UNSTABLE_STALE_DURATION: 5 minutes

DEFAULT_NS_POLLING: 3 times

9. IANA Considerations

This document has no request to IANA.

10. Future Work

Future documents may extend this specification by allowing the 6BBR to redistribute host routes in routing protocols that would operate over the backbone, or in MIPv6, or FMIP, or the Locator/ID Separation Protocol (LISP) [RFC6830] to support mobility on behalf of the 6LNs, etc...

11. Acknowledgments

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

12. References

12.1. Normative References

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

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

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <https://www.rfc-editor.org/info/rfc4291>.

[RFC4429] Moore, N., "Optimistic Duplicate Address Detection (DAD) for IPv6", RFC 4429, DOI 10.17487/RFC4429, April 2006, <https://www.rfc-editor.org/info/rfc4429>.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <https://www.rfc-editor.org/info/rfc4861>.

[RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <https://www.rfc-editor.org/info/rfc4862>.

[RFC6059] Krishnan, S. and G. Daley, "Simple Procedures for Detecting Network Attachment in IPv6", RFC 6059, DOI 10.17487/RFC6059, November 2010, <https://www.rfc-editor.org/info/rfc6059>.

[RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <https://www.rfc-editor.org/info/rfc6775>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

12.2. Informative References

[I-D.ietf-6lo-ap-nd]
Thubert, P., Sarikaya, B., Sethi, M., and R. Struik, "Address Protected Neighbor Discovery for Low-power and Lossy Networks", draft-ietf-6lo-ap-nd-08 (work in progress), October 2018.

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

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

[I-D.ietf-mboned-ieee802-mcast-problems]
Perkins, C., McBride, M., Stanley, D., Kumari, W., and J. Zuniga, "Multicast Considerations over IEEE 802 Wireless Media", draft-ietf-mboned-ieee802-mcast-problems-02 (work in progress), August 2018.

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

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

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

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

[RFC4389] Thaler, D., Talwar, M., and C. Patel, "Neighbor Discovery Proxies (ND Proxy)", RFC 4389, DOI 10.17487/RFC4389, April 2006, <https://www.rfc-editor.org/info/rfc4389>.

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

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

[RFC5415] Calhoun, P., Ed., Montemurro, M., Ed., and D. Stanley, Ed., "Control And Provisioning of Wireless Access Points (CAPWAP) Protocol Specification", RFC 5415, DOI 10.17487/RFC5415, March 2009, <https://www.rfc-editor.org/info/rfc5415>.

[RFC6275] Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, DOI 10.17487/RFC6275, July 2011, <https://www.rfc-editor.org/info/rfc6275>.

[RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", RFC 6830, DOI 10.17487/RFC6830, January 2013, <https://www.rfc-editor.org/info/rfc6830>.

[RFC7048] Nordmark, E. and I. Gashinsky, "Neighbor Unreachability Detection Is Too Impatient", RFC 7048, DOI 10.17487/RFC7048, January 2014, <https://www.rfc-editor.org/info/rfc7048>.

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

[RFC7559] Krishnan, S., Anipko, D., and D. Thaler, "Packet-Loss Resiliency for Router Solicitations", RFC 7559, DOI 10.17487/RFC7559, May 2015, <https://www.rfc-editor.org/info/rfc7559>.

[RFC7772] Yourtchenko, A. and L. Colitti, "Reducing Energy Consumption of Router Advertisements", BCP 202, RFC 7772, DOI 10.17487/RFC7772, February 2016, <https://www.rfc-editor.org/info/rfc7772>.

12.3. External Informative References

[IEEEstd8021] IEEE standard for Information Technology, "IEEE Standard for Information technology -- Telecommunications and information exchange between systems Local and metropolitan area networks Part 1: Bridging and Architecture".

[IEEEstd80211]
IEEE standard for Information Technology, "IEEE Standard for Information technology -- Telecommunications and information exchange between systems Local and metropolitan area networks-- Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications".

[IEEEstd802151]
IEEE standard for Information Technology, "IEEE Standard for Information Technology - Telecommunications and Information Exchange Between Systems - Local and Metropolitan Area Networks - Specific Requirements. - Part 15.1: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Wireless Personal Area Networks (WPANs)".

[IEEEstd802154]
IEEE standard for Information Technology, "IEEE Standard for Local and metropolitan area networks -- Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs)".

Appendix A. Changes from revision 07 to revision 08

This section lists the changes between draft-ietf-6lo-backbone-router revisions ...-07.txt and ...-08.txt.

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

Authors' Addresses

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

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

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

Email: charliep@computer.org

6lo
Internet-Draft
Intended status: Standards Track
Expires: April 18, 2019

Lijo Thomas
C-DAC
S. Anamalamudi
SRM University-AP
S.V.R.Anand
Malati Hegde
Indian Institute of Science
C. Perkins
Futurewei
October 15, 2018

Packet Delivery Deadline time in 6LoWPAN Routing Header
draft-ietf-6lo-deadline-time-03

Abstract

This document specifies a new type for the 6LoWPAN routing header containing the delivery deadline time for data packets. The deadline time enables forwarding and scheduling decisions for time critical IoT M2M applications that need deterministic delay guarantees over constrained networks and operate within time-synchronized networks.

Status of This Memo

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

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

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

This Internet-Draft will expire on April 18, 2019.

Copyright Notice

Table of Contents

1. Introduction

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

The draft [I-D.ietf-roll-routing-dispatch] specifies the 6LoWPAN Routing Header (6LoRH), compression schemes for RPL routing (source routing) operation [RFC6554], header compression of RPL Packet Information [RFC6553], and IP-in-IP encapsulation. This document specifies a new Deadline-6LoRHE type for the 6LoWPAN Dispatch Page 1, so that the deadline time of data packets can be included within the 6LoWPAN routing header. This document also specifies handling of the

deadline time when packets traverse through time-synchronized networks operating in different timezones or distinct reference clocks. Time synchronization techniques need not be mandated by this specificiation. There are a number of standards available for this purpose, including IEEE 1588 [ieee-1588], IEEE 802.1AS [dot1AS-2011], IEEE 802.15.4-2015 TSCH [dot15-tsch], and more.

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

2. Terminology

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

This document uses terminology consistent with the terminology used in [RFC6550] and [I-D.ietf-6tisch-terminology]. Also, in this document, the terms "expiration time", "delivery deadline time", and "deadline" are used interchangeably with the same meaning.

3. 6LoRHE Generic Format

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

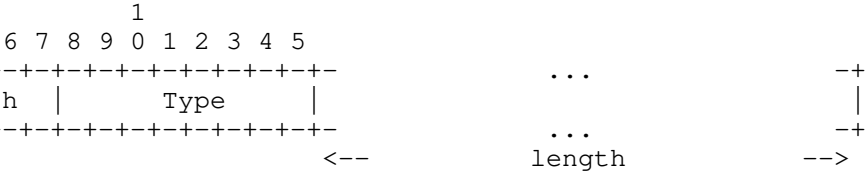

Figure 1: 6LoRHE format

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

- o Type: Type of the 6LoRHE.

- o length: variable

4. Deadline-6LoRHE

The Deadline-6LoRHE (see Figure 3) is an elective 6LoRH (i.e., a 6LoRHE [RFC8138]) that provides the Deadline Time (DT) for an IPv6 datagram in a compressed form. Along with the deadline, the header can include the packet Origination Time (OT), the time at which the packet is enqueued for transmission, to enable a close estimate of the total delay incurred by a packet. The OT field is initialized by the sender using the current time at the outgoing network interface through which the packet is forwarded.

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

```
packet_deadline_time = packet_origination_time + max_delay
```

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

The delay experienced by packets in the network is a useful metric for network diagnostics and performance monitoring. Whenever the packets crosses into a network using a different reference clock, the Origination Time field is updated to represent the same Origination Time, but expressed using the reference clock of the interface into the new network. This is the same as the current time when the packet is transmitted into the new network, minus the delay already experienced by the packet, say 't'. In this way, within the newly entered network, the packet will appear to have originated 't' time units earlier with respect to the reference clock of the new network.

```
Origination Time in new network = current_time_in_new_network -
                delay_already_experienced_in_previous_network(s)
```

The following example illustrates the origination time calculation when a packet travels between three networks, each in a different time zone. 'x' can be 1,2 or 3.

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

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

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

TZx : Indicates the time zone of network 'x'

As an illustration, we consider a packet traversing through three time synchronized networks along with numerical values as shown in Figure 1.

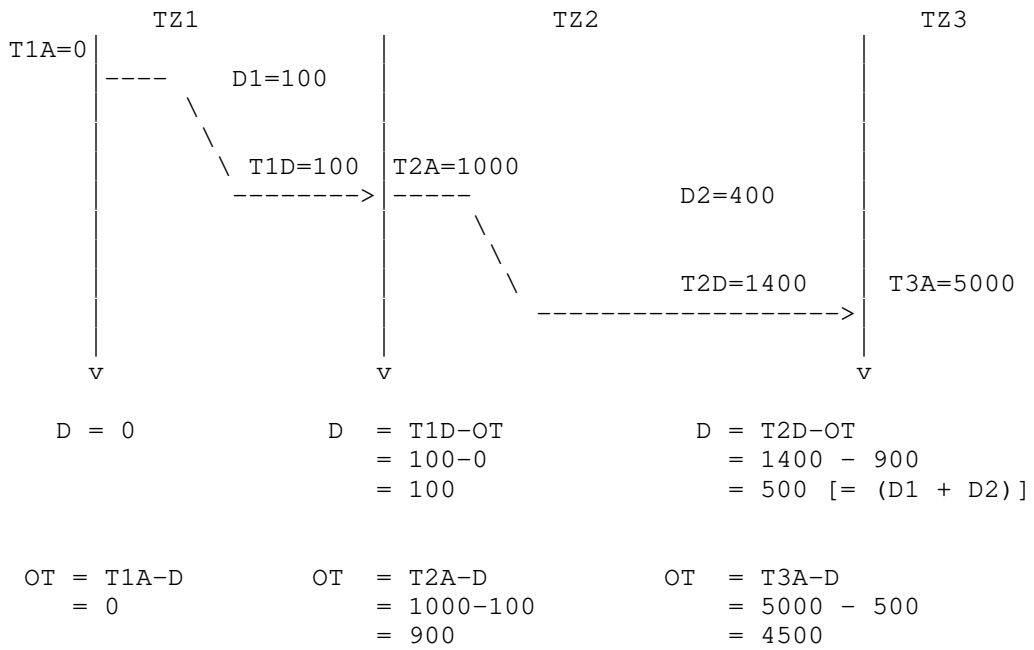

Figure 2: Origination Time update example

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

5. Deadline-6LoRHE Format

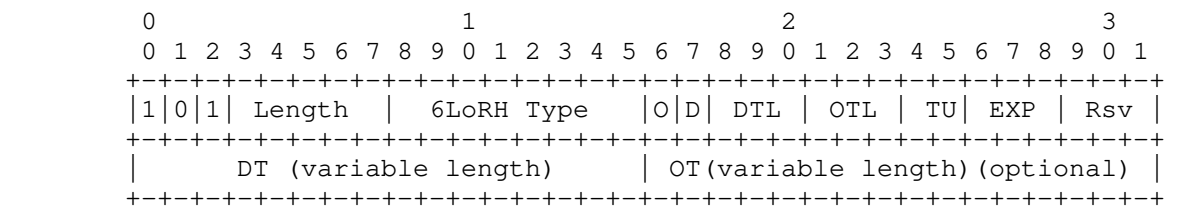

Figure 3: Deadline-6LoRHE format

Length (5 bits): Length represents the total length of the Deadline-6LoRHE type measured in octets.

6LoRH Type: TBD

O flag (1bit): Indicates the presence of Origination Time field. '1' means the OT field is present, and '0' means it is absent.

D flag (1 bit): The 'D' flag, set by the Sender, indicates the action to be taken when a 6LR detects that the deadline time has elapsed. If 'D' bit is 1, then the 6LR MUST drop the packet if the deadline time is elapsed.

If 'D' bit is 0, implies the packet MAY be forwarded on an exception basis, if the forwarding node is NOT in a situation of constrained resource, and if there are reasons to suspect that downstream nodes might find it useful (delay measurements, interpolations, etc.).

DTL (3 bits): Length of DT field as an unsigned 3-bit integer, encoding the length of the field in octets, minus one.

OTL (3 bits) : Length of OT field as an unsigned 3-bit integer, encoding the length of the field in octets, minus one.

For example, DTL = 000 means the deadline time in the 6LoRHE is 1 octet (8 bits) long. Similarly, OTL = 111 means the origination time is 8 octets (64 bits) long.

TU (2 bits) : Indicates the time units for DT and OT fields

00 : Time represented in microseconds
01 : Time represented in seconds
10 : Network ASN
11 : Reserved

EXP (3 bits) : Multiplication factor expressed as exponent of 10.

The value of the DT field is multiplied by 10 to this power, to get the actual deadline time in the units represented by TU. The default value of EXP is 000, so that the DT field is unaffected.

Rsv (3 bits) : Reserved, sent as zero and ignored on receipt

DT Value (8..64-bit) : An unsigned integer of DTL octets giving the Deadline Time value

OT Value (8..64-bit) : An unsigned integer of OTL octets giving the Origination Time value

Whenever a sender initiates the IP datagram, it includes the Deadline-6LoRHE along with other 6LoRH information.

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

deadline_time = packet_origination_time + max_delay
= 55400 + 100 (in Network ASNs)
= 55500(Network ASNs)

Deadline-6LoRHE encoding with 'O' flag and 'D' flag set to 1:

DTL = 001, OTL = 001, TU = '10', EXP = 2, DT = 0x22B, OT = 0x22A

6. Deadline-6LoRHE in Three Network Scenarios

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

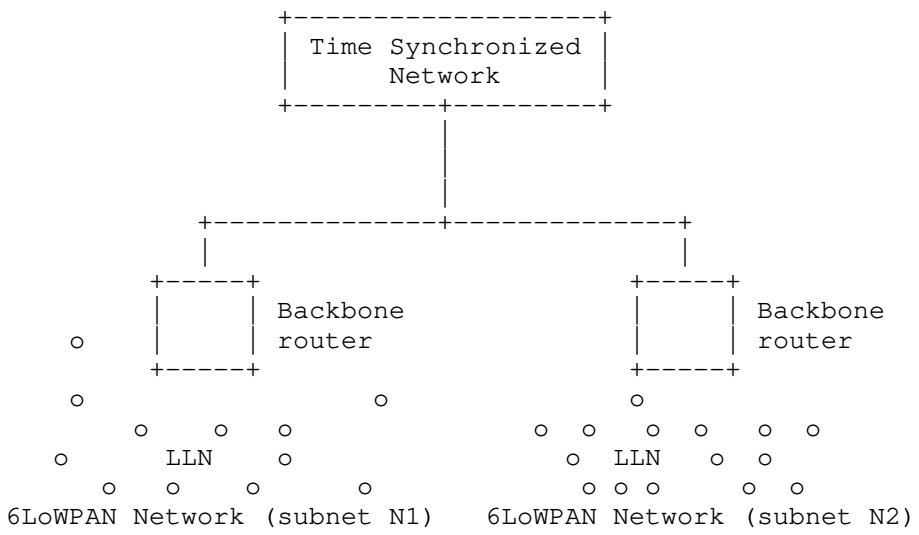

Figure 4: Intra-network Timezone Scenario

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

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

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

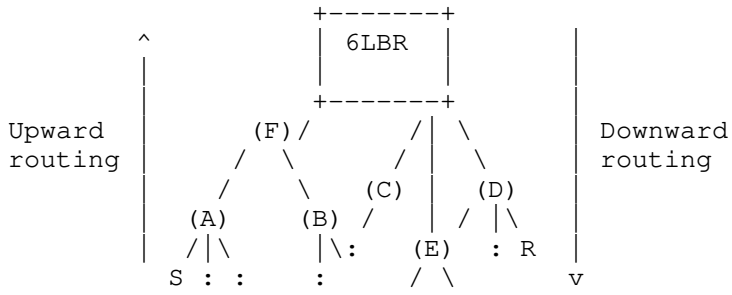

Figure 5: End points within same DODAG (subnet N1)

At the tunnel endpoint of the IPv6-in-IPv6 encapsulation, the Deadline-6LoRHE is copied back from the outer header to inner header, and the inner IP packet is delivered to 'R'.

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

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

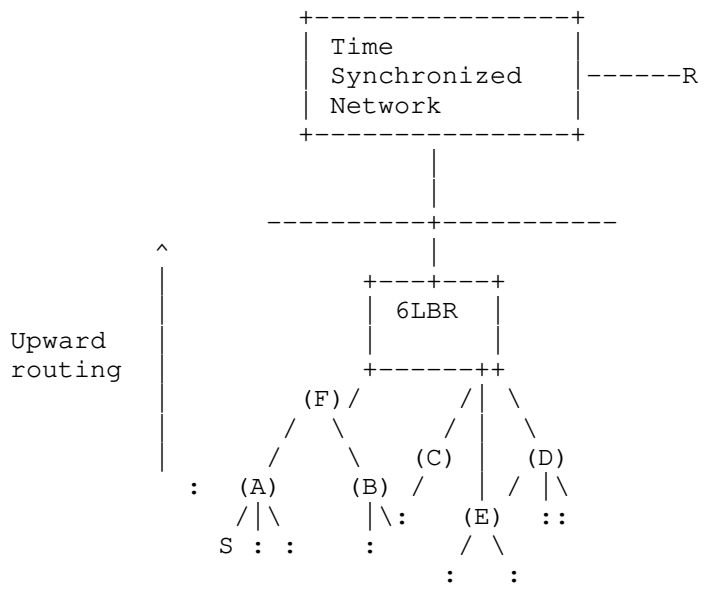

Figure 6: Packet transmission in Dissimilar L2 Technologies or Internet

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

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

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

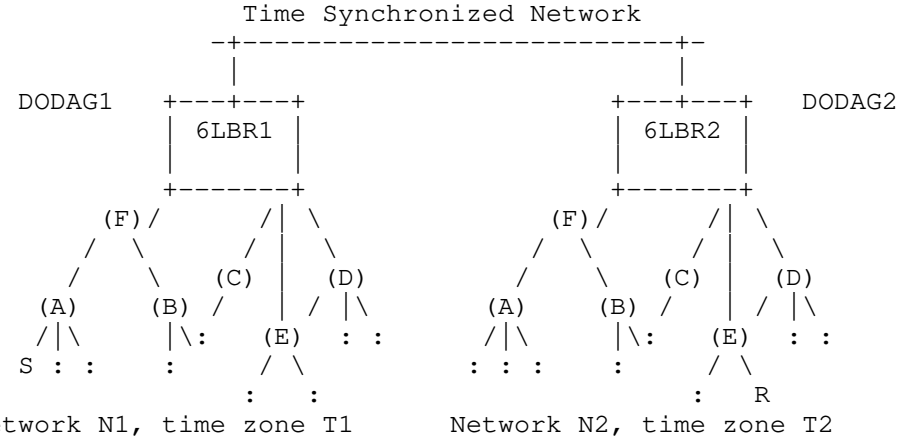

Figure 7: Packet transmission in different DODAGs(N1 to N2)

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

```
current_time = ASN at LBR * slot_length_value

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

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

7. IANA Considerations

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

Elective 6LoRH Type	Value
Deadline-6LoRHE	TBD

Figure 8: Deadline-6LoRHE type

8. Security Considerations

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

9. Acknowledgements

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

10. References

10.1. Normative References

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

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

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

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

[RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <https://www.rfc-editor.org/info/rfc6282>.

[RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <https://www.rfc-editor.org/info/rfc6550>.

[RFC6553] Hui, J. and JP. Vasseur, "The Routing Protocol for Low-Power and Lossy Networks (RPL) Option for Carrying RPL Information in Data-Plane Datagrams", RFC 6553, DOI 10.17487/RFC6553, March 2012, <https://www.rfc-editor.org/info/rfc6553>.

[RFC6554] Hui, J., Vasseur, JP., Culler, D., and V. Manral, "An IPv6 Routing Header for Source Routes with the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6554, DOI 10.17487/RFC6554, March 2012, <https://www.rfc-editor.org/info/rfc6554>.

[RFC8138] Thubert, P., Ed., Bormann, C., Toutain, L., and R. Cragie, "IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing Header", RFC 8138, DOI 10.17487/RFC8138, April 2017, <https://www.rfc-editor.org/info/rfc8138>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

10.2. Informative References

[dot15-tsch]
P802.11, "IEEE Standard for Low-Rate Wireless Networks, Part 15.4, IEEE Std 802.15.4-2015", April 2016.

[dot1AS-2011]
IEEE 802.1AS Working Group, "IEEE Standard for Local and Metropolitan Area Networks - Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks", March 2011.

[dotBLEMesh]
Luca Leonardi, Gaetano Pattim, and Lucia Lo Bello, "Multi-Hop Real-Time Communications Over Bluetooth Low Energy Industrial Wireless Mesh Networks", IEEE Access Vol 6, 26505-26519, May 2018.

[dotWi-SUN]
Hiroshi Harada, Keiichi Mizutani, Jun Fujiwara, Kentaro Mochizuki, Kentaro Obata, and Okumura, Ryota, "IEEE 802.15.4g Based Wi-SUN Communication Systems", IEICE Transactions on Communications volume E100.B, Jan 2017.

[I-D.ietf-6lo-backbone-router]
Thubert, P. and C. Perkins, "IPv6 Backbone Router", draft-ietf-6lo-backbone-router-07 (work in progress), September 2018.

[I-D.ietf-6lo-blemesh]
Gomez, C., Darroudi, S., Savolainen, T., and M. Spoerk, "IPv6 Mesh over BLUETOOTH(R) Low Energy using IPSP", draft-ietf-6lo-blemesh-03 (work in progress), July 2018.

[I-D.ietf-detnet-use-cases]
Grossman, E., "Deterministic Networking Use Cases", draft-ietf-detnet-use-cases-19 (work in progress), October 2018.

[I-D.ietf-ippm-ioam-data]
Brockners, F., Bhandari, S., Pignataro, C., Gredler, H., Leddy, J., Youell, S., Mizrahi, T., Mozes, D., Lapukhov, P., Chang, R., daniel.bernier@bell.ca, d., and J. Lemon, "Data Fields for In-situ OAM", draft-ietf-ippm-ioam-data-03 (work in progress), June 2018.

[I-D.ietf-roll-useofrplinfo]
Robles, I., Richardson, M., and P. Thubert, "When to use RFC 6553, 6554 and IPv6-in-IPv6", draft-ietf-roll-useofrplinfo-23 (work in progress), May 2018.

[ieee-1588]
Precise Time and Time Interval Working Group, "IEEE Std 1588-2008 Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", July 2008.

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

Appendix A. Changes from revision 02 to revision 03

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

- Added non-normative 6LoRHE description, citing RFC 8138.
- Specified that the Origination Time (OT) is the time that packet is enqueued for transmission.
- Mentioned more sources of packet delay.
- Clarified reasons that packet MAY be forwarded if 'D' bit is 0.
- Clarified that DT, OT, DTL and OTL are unsigned integers.
- Updated bibliographic citations, including BLEmesh and Wi-SUN.

Appendix B. Changes from revision 01 to revision 02

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

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

Appendix C. Changes between earlier versions

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

- Changed "SHOULD drop" to "MUST drop" a packet if the deadline is passed (see Section 5).
- Added explanatory text about how packet delays might arise. (see Section 4).

- Mentioned availability of time-synchronization protocols (see Section 1).
- Updated bibliographic citations.
- Alphabetized contributor names.
- Added this section.

Authors' Addresses

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

Email: lijo@cdac.in

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

Email: satishnaidu80@gmail.com

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

Email: anand@ece.iisc.ernet.in

Malati Hegde
Indian Institute of Science
Bangalore 560012
India

Email: malati@ece.iisc.ernet.in

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

Email: charliep@computer.org

6lo
Internet-Draft
Updates: 4944 (if approved)
Intended status: Standards Track
Expires: March 29, 2019

P. Thubert, Ed.
Cisco Systems

September 20, 2018

6LoWPAN Selective Fragment Recovery
draft-ietf-6lo-fragment-recovery-00

Abstract

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

Status of This Memo

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

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

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

This Internet-Draft will expire on December 29, 2018.

Table of Contents

1. Introduction

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

"Transmission of IPv6 Packets over IEEE 802.15.4 Networks" [RFC4944] defines the original 6LoWPAN datagram fragmentation mechanism for

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

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

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

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

2. Updating RFC 4944

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

3. Updating draft-watteyne-6lo-minimal-fragment

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

At the time of this writing, [I-D.watteyne-6lo-minimal-fragment] allows for refragmenting in intermediate nodes, meaning that some bytes from a given fragment may be left in the VRB to be added to the next fragment. The reason for this to happen would be the need for space in the outgoing fragment that was not needed in the incoming fragment, for instance because the 6LoWPAN Header Compression is not as efficient on the outgoing link, e.g., if the IID of the source IPv6 address is elided on the first hop because it matches the MAC address, but cannot be on the next hops. This specification does not allow this since fragments are recovered end-to-end. This means that the fragments that contain 6LoWPAN-compressed data must have enough slack in them to enable a lesser efficient compression in the next hops to still fit in one frame.

4. Terminology

4.1. 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.

4.2. References

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

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

4.3. 6LoWPAN Acronyms

This document uses the following acronyms:

6BBR: 6LoWPAN Backbone Router

6LBR: 6LoWPAN Border Router

6LN: 6LoWPAN Node

6LR: 6LoWPAN Router

LLN: Low-Power and Lossy Network

4.4. Referenced Work

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

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

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

Readers are expected to be familiar with all the terms and concepts that are discussed in "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals" [RFC4919] and "Transmission of IPv6 Packets over IEEE 802.15.4 Networks" [RFC4944].

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

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

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

4.5. New Terms

This specification uses the following terms:

6LoWPAN endpoints

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

5. New Dispatch types and headers

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

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

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

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

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

Figure 1: Additional Dispatch Value Bit Patterns

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

5.1. Recoverable Fragment Dispatch type and Header

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

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

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

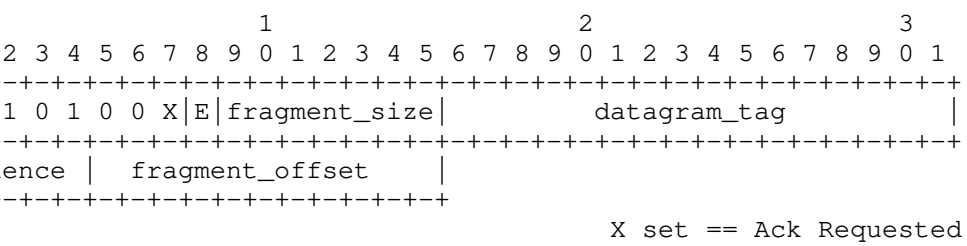

Figure 2: RFRAG Dispatch type and Header

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

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

Fragment_size: 7 bit unsigned integer; the size of this fragment in a unit that depends on the MAC layer technology. For IEEE Std. 802.15.4, the unit is octet, and the maximum fragment size, which is constrained by the maximum frame size of 128 octet minus the overheads of the MAC and Fragment Headers, is not limited by this encoding.

Sequence: 5 bit unsigned integer; the sequence number of the fragment. Fragments are sequence numbered [0..N] where N is in [0..31]. A sequence of 0 indicates the first fragment in a datagram. For IEEE Std. 802.15.4, as long as the overheads enable a fragment size of 64 octets or more, this enables to fragment a packet of 2047 octets.

Fragment_offset: 11 bit unsigned integer;

* When set to a non-0 value, the semantics of the Fragment_offset depends on the value of the Sequence.

  + When the Sequence is not 0, this field indicates the offset of the fragment in the compressed form. The fragment should be forwarded based on an existing VRB as described in Section 7.2, or silently dropped if none is found.

  + For a first fragment (i.e. with a sequence of 0), this field is overloaded to indicate the total_size of the compressed packet, to help the receiver allocate an adapted buffer for the reception and reassembly operations. This format limits the maximum MTU on a 6LoWPAN link to 2047 bytes, but 1280 bytes is the recommended value to avoid issues with IPV6 Path MTU Discovery [RFC8201]. The fragment should be routed based on the destination IPv6 address, and an VRB state should be installed as described in Section 7.1.

* When set to 0, this field indicates an abort condition and all state regarding the datagram should be cleaned up once the processing of the fragment is complete; the processing of the fragment depends on whether there is a VRB already established for this datagram, and the next hop is still reachable:

  + if a VRB already exists and is not broken, the fragment is to be forwarded along the associated Label Switched Path (LSP) as described in Section 7.2, but regardless of the value of the Sequence field;

  + else, if the Sequence is 0, then the fragment is to be routed as described in Section 7.1 but no state is conserved afterwards.

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

5.2. RFRAG Acknowledgment Dispatch type and Header

This specification also defines a 4-octet RFRAG Acknowledgment bitmap that is used by the reassembling end point to confirm selectively the reception of individual fragments. A given offset in the bitmap maps one to one with a given sequence number.

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

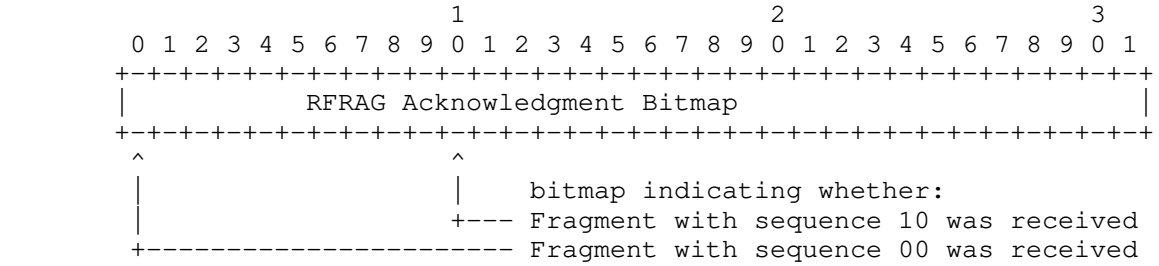

Figure 3: RFRAG Acknowledgment bitmap encoding

Figure 4 shows an example Acknowledgment bitmap which indicates that all fragments from sequence 0 to 20 were received, except for fragments 1, 2 and 16 that were either lost or are still in the network over a slower path.

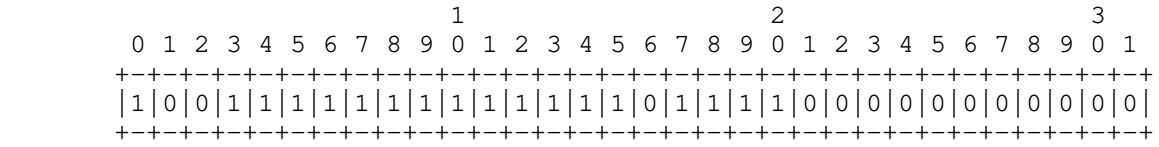

Figure 4: Expanding 3 octets encoding

The RFRAG Acknowledgment Bitmap is included in a RFRAG Acknowledgment header, as follows:

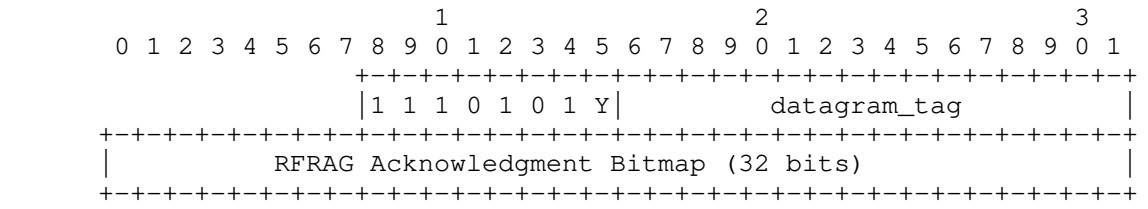

Figure 5: RFRAG Acknowledgment Dispatch type and Header

Y: 1 bit; Explicit Congestion Notification Echo

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

RFRAG Acknowledgment Bitmap

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

6. Fragments Recovery

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

The 6LoWPAN endpoint that fragments the packets at 6LoWPAN level (the sender) also controls when the reassembling end point sends the RFRAG Acknowledgments by setting the Ack Requested flag in the RFRAG packets. It may set the Ack Requested flag on any fragment to perform congestion control by limiting the number of outstanding fragments, which are the fragments that have been sent but for which reception or loss was not positively confirmed by the reassembling endpoint. When the sender of the fragment knows that an underlying link-layer mechanism protects the Fragments, it may refrain from

using the RFRAG Acknowledgment mechanism, and never set the Ack Requested bit. When it receives a fragment with the ACK Request flag set, the 6LoWPAN endpoint that reassembles the packets at 6LoWPAN level (the receiver) sends back an RFRAG Acknowledgment to confirm reception of all the fragments it has received so far.

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

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

An observation is that streamlining forwarding of fragments generally reduces the latency over the LLN mesh, providing room for retries within existing upper-layer reliability mechanisms. The sender protects the transmission over the LLN mesh with a retry timer that is computed according to the method detailed in [RFC6298]. It is expected that the upper layer retries obey the recommendations in "UDP Usage Guidelines" [RFC8085], in which case a single round of fragment recovery should fit within the upper layer recovery timers.

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

When a single frequency is used by contiguous hops, the sender should wait a reasonable amount of time between fragments so as to let a fragment progress a few hops and avoid hidden terminal issues. This precaution is not required on channel hopping technologies such as Time Slotted CHannel Hopping (TSCH) [RFC6554]

When the sender decides that a packet should be dropped and the fragmentation process canceled, it sends a pseudo fragment with the

fragment_offset, sequence and fragment_size all set to 0, and no data. Upon reception of this message, the receiver should clean up all resources for the packet associated to the datagram_tag. If an acknowledgment is requested, the receiver responds with a NULL bitmap.

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

7. Forwarding Fragments

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

This specification extends the Virtual Reassembly Buffer (VRB) technique to forward fragments with no intermediate reconstruction of the entire packet. The first fragment carries the IP header and it is routed all the way from the fragmenting end point to the reassembling end point. Upon the first fragment, the routers along the path install a label-switched path (LSP), and the following fragments are label-switched along that path. As a consequence, alternate routes not possible for individual fragments. The datagram_tag is used to carry the label, that is swapped at each hop. All fragments follow the same path and fragments are delivered in the order at which they are sent.

7.1. Upon the first fragment

In Route-Over mode, the source and destination MAC addressed in a frame change at each hop. The label that is formed and placed in the datagram_tag is associated to the source MAC and only valid (and unique) for that source MAC. Upon a first fragment (i.e. with a sequence of zero), a VRB and the associated LSP state are created for the tuple (source MAC address, datagram_tag) and the fragment is forwarded along the IPv6 route that matches the destination IPv6 address in the IPv6 header as prescribed by [I-D.watteyne-6lo-minimal-fragment]. The LSP state enables to match the (previous MAC address, datagram_tag) in an incoming fragment to the tuple (next MAC address, swapped datagram_tag) used in the forwarded fragment and points at the VRB. In addition, the router also forms a Reverse LSP state indexed by the MAC address of the next

hop and the swapped datagram_tag. This reverse LSP state also points at the VRB and enables to match the (next MAC address, swapped_datagram_tag) found in an RFRAG Acknowledgment to the tuple (previous MAC address, datagram_tag) used when forwarding a Fragment Acknowledgment (RFRAG-ACK) back to the sender endpoint.

7.2. Upon the next fragments

Upon a next fragment (i.e. with a non-zero sequence), the router looks up a LSP indexed by the tuple (MAC address, datagram_tag) found in the fragment. If it is found, the router forwards the fragment using the associated VRB as prescribed by [I-D.watteyne-6lo-minimal-fragment].

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

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

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

7.3. Upon the RFRAG Acknowledgments

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

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

Either way, if the RFRAG-ACK indicates either an error (NULL bitmap) or that the fragment was entirely received (FULL bitmap), arms a short timer, and upon timeout, the VRB and all associate state are destroyed. During that time, fragments of that datagram may still be received, e.g. if the RFRAG-ACK was lost on the way back and the

source retried the last fragment. In that case, the router sends an abort RFRAG-ACK along the Reverse LSP to complete the clean up.

8. Security Considerations

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

9. IANA Considerations

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

10. Acknowledgments

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

11. References

11.1. Normative References

[I-D.watteyne-6lo-minimal-fragment]
Watteyne, T., Bormann, C., and P. Thubert, "LLN Minimal Fragment Forwarding", draft-watteyne-6lo-minimal-fragment-01 (work in progress), March 2018.

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

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

[RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <https://www.rfc-editor.org/info/rfc6282>.

[RFC6554] Hui, J., Vasseur, JP., Culler, D., and V. Manral, "An IPv6 Routing Header for Source Routes with the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6554, DOI 10.17487/RFC6554, March 2012, <https://www.rfc-editor.org/info/rfc6554>.

[RFC8025] Thubert, P., Ed. and R. Cragie, "IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Paging Dispatch", RFC 8025, DOI 10.17487/RFC8025, November 2016, <https://www.rfc-editor.org/info/rfc8025>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

11.2. Informative References

[I-D.ietf-6tisch-architecture]
Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", draft-ietf-6tisch-architecture-14 (work in progress), April 2018.

[IEEE.802.15.4]
IEEE, "IEEE Standard for Low-Rate Wireless Networks", IEEE Standard 802.15.4, DOI 10.1109/IEEE P802.15.4-REVd/D01, <http://ieeexplore.ieee.org/document/7460875/>.

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

[RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, DOI 10.17487/RFC3031, January 2001, <https://www.rfc-editor.org/info/rfc3031>.

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

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

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

[RFC5681] Allman, M., Paxson, V., and E. Blanton, "TCP Congestion Control", RFC 5681, DOI 10.17487/RFC5681, September 2009, <https://www.rfc-editor.org/info/rfc5681>.

[RFC6298] Paxson, V., Allman, M., Chu, J., and M. Sargent, "Computing TCP's Retransmission Timer", RFC 6298, DOI 10.17487/RFC6298, June 2011, <https://www.rfc-editor.org/info/rfc6298>.

[RFC6606] Kim, E., Kaspar, D., Gomez, C., and C. Bormann, "Problem Statement and Requirements for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing", RFC 6606, DOI 10.17487/RFC6606, May 2012, <https://www.rfc-editor.org/info/rfc6606>.

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

[RFC7567] Baker, F., Ed. and G. Fairhurst, Ed., "IETF Recommendations Regarding Active Queue Management", BCP 197, RFC 7567, DOI 10.17487/RFC7567, July 2015, <https://www.rfc-editor.org/info/rfc7567>.

[RFC8085] Eggert, L., Fairhurst, G., and G. Shepherd, "UDP Usage Guidelines", BCP 145, RFC 8085, DOI 10.17487/RFC8085, March 2017, <https://www.rfc-editor.org/info/rfc8085>.

[RFC8087] Fairhurst, G. and M. Welzl, "The Benefits of Using Explicit Congestion Notification (ECN)", RFC 8087, DOI 10.17487/RFC8087, March 2017, <https://www.rfc-editor.org/info/rfc8087>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

[RFC8201] McCann, J., Deering, S., Mogul, J., and R. Hinden, Ed., "Path MTU Discovery for IP version 6", STD 87, RFC 8201, DOI 10.17487/RFC8201, July 2017, <https://www.rfc-editor.org/info/rfc8201>.

Appendix A. Rationale

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

The list of those usages includes:

Towards the LLN node:

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

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

From the LLN node:

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

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

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

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

When a fragment is lost in transmission, the lack of recovery in the original fragmentation system of RFC 4944 implies that all fragments

are resent, further contributing to the congestion that caused the initial loss, and potentially leading to congestion collapse.

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

Considering that RFC 4944 defines an MTU is 1280 bytes and that in most incarnations (but 802.15.4g) a IEEE Std. 802.15.4 frame can limit the MAC payload to as few as 74 bytes, a packet might be fragmented into at least 18 fragments at the 6LoWPAN shim layer. Taking into account the worst-case header overhead for 6LoWPAN Fragmentation and Mesh Addressing headers will increase the number of required fragments to around 32. This level of fragmentation is much higher than that traditionally experienced over the Internet with IPv4 fragments. At the same time, the use of radios increases the probability of transmission loss and Mesh-Under techniques compound that risk over multiple hops.

Mechanisms such as TCP or application-layer segmentation could be used to support end-to-end reliable transport. One option to support bulk data transfer over a frame-size-constrained LLN is to set the Maximum Segment Size to fit within the link maximum frame size. Doing so, however, can add significant header overhead to each 802.15.4 frame. In addition, deploying such a mechanism requires that the end-to-end transport is aware of the delivery properties of the underlying LLN, which is a layer violation, and difficult to achieve from the far end of the IPv6 network.

Appendix B. Requirements

For one-hop communications, a number of Low Power and Lossy Network (LLN) link-layers propose a local acknowledgment mechanism that is enough to detect and recover the loss of fragments. In a multihop environment, an end-to-end fragment recovery mechanism might be a good complement to a hop-by-hop MAC level recovery. This draft introduces a simple protocol to recover individual fragments between 6LoWPAN endpoints that may be multiple hops away. The method addresses the following requirements of a LLN:

Number of fragments

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

Minimum acknowledgment overhead

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

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

Controlled latency

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

Optional congestion control

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

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

Appendix C. Considerations On Flow Control

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

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

Congestion on the forward path can also be indicated by an Explicit Congestion Notification (ECN) mechanism. Though whether and how ECN [RFC3168] is carried out over the LoWPAN is out of scope, this draft provides a way for the destination endpoint to echo an ECN indication back to the source endpoint in an acknowledgment message as represented in Figure 5 in Section 5.2.

It must be noted that congestion and collision are different topics. In particular, when a mesh operates on a same channel over multiple hops, then the forwarding of a fragment over a certain hop may

collide with the forwarding of a next fragment that is following over a previous hop but in a same interference domain. This draft enables an end-to-end flow control, but leaves it to the sender stack to pace individual fragments within a transmit window, so that a given fragment is sent only when the previous fragment has had a chance to progress beyond the interference domain of this hop. In the case of 6TiSCH [I-D.ietf-6tisch-architecture], which operates over the TimeSlotted Channel Hopping [RFC7554] (TSCH) mode of operation of IEEE802.14.5, a fragment is forwarded over a different channel at a different time and it makes full sense to transmit the next fragment as soon as the previous fragment has had its chance to be forwarded at the next hop.

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

From the sender standpoint, all outstanding fragments might still be in the network and contribute to its congestion. There is an assumption, though, that after a certain amount of time, a frame is either received or lost, so it is not causing congestion anymore. This amount of time can be estimated based on the round trip delay between the 6LoWPAN endpoints. The method detailed in [RFC6298] is recommended for that computation.

The reader is encouraged to read through "Congestion Control Principles" [RFC2914]. Additionally [RFC7567] and [RFC5681] provide deeper information on why this mechanism is needed and how TCP handles Congestion Control. Basically, the goal here is to manage the amount of fragments present in the network; this is achieved by to reducing the number of outstanding fragments over a congested path by throttling the sources.

Section 6 describes how the sender decides how many fragments are (re)sent before an acknowledgment is required, and how the sender adapts that number to the network conditions.

Author’s Address

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

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

6Lo Working Group
Internet-Draft
Intended status: Standards Track
Expires: February 24, 2021

Y. Choi, Ed.
Y-G. Hong
ETRI
J-S. Youn
Dongeui Univ
D-K. Kim
KNU
J-H. Choi
Samsung Electronics Co.,
August 23, 2020

Transmission of IPv6 Packets over Near Field Communication
draft-ietf-6lo-nfc-17

Abstract

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

Status of This Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at 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 February 24, 2021.

Table of Contents

1. Introduction

NFC is a set of short-range wireless technologies, typically requiring a distance between sender and receiver of 10 cm or less. NFC operates at 13.56 MHz, and at rates ranging from 106 kbit/s to 424 kbit/s, as per the ISO/IEC 18000-3 air interface [ECMA-340]. NFC

builds upon RFID systems by allowing two-way communication between endpoints. NFC always involves an initiator and a target; the initiator actively generates an RF field that can power a passive target. This enables NFC targets to take very simple form factors, such as tags, stickers, key fobs, or cards, while avoiding the need for batteries. NFC peer-to-peer communication is possible, provided that both devices are powered. As of the writing, NFC is supported by the main smartphone operating systems.

NFC is often regarded as a secure communications technology, due to its very short transmission range.

In order to benefit from Internet connectivity, it is desirable for NFC-enabled devices to support IPv6, considering its large address space, along with tools for unattended operation, among other advantages. This document specifies how IPv6 is supported over NFC by using IPv6 over Low-power Wireless Personal Area Network (6LoWPAN) techniques [RFC4944], [RFC6282], [RFC6775]. 6LoWPAN is suitable, considering that it was designed to support IPv6 over IEEE 802.15.4 networks, and some of the characteristics of the latter are similar to those of NFC.

2. Conventions and Terminology

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

3. Overview of Near Field Communication Technology

This section presents an overview of NFC, focusing on the characteristics of NFC that are most relevant for supporting IPv6.

NFC enables simple, two-way, interaction between two devices, allowing users to perform contactless transactions, access digital content, and connect electronic devices with a single touch. NFC utilizes key elements in existing standards for contactless card Technology, such as ISO/IEC 14443 A&B and JIS-X 6319-4. NFC allows devices to share information at a distance up to 10 cm with a maximum physical layer bit rate of 424 kbps.

3.1. Peer-to-peer Mode of NFC

NFC defines three modes of operation: card emulation, peer-to-peer, and reader/writer. Only the peer-to-peer mode allows two NFC-enabled devices to communicate with each other to exchange information

bidirectionally. The other two modes do not support two-way communications between two devices. Therefore, the peer-to-peer mode is used for IPv6 over NFC.

3.2. Protocol Stack of NFC

NFC defines a protocol stack for the peer-to-peer mode (Figure 1). The peer-to-peer mode is offered by the Activities Digital Protocol at the NFC Physical Layer. The NFC Logical Link Layer comprises the Logical Link Control Protocol (LLCP), and when IPv6 is used over NFC, it also includes an IPv6-LLCP Binding. IPv6 and its underlying adaptation Layer (i.e., IPv6-over-NFC adaptation layer) are placed directly on the top of the IPv6-LLCP Binding. An IPv6 datagram is transmitted by the Logical Link Control Protocol (LLCP) with reliable, two-way transmission of information between the peer devices.

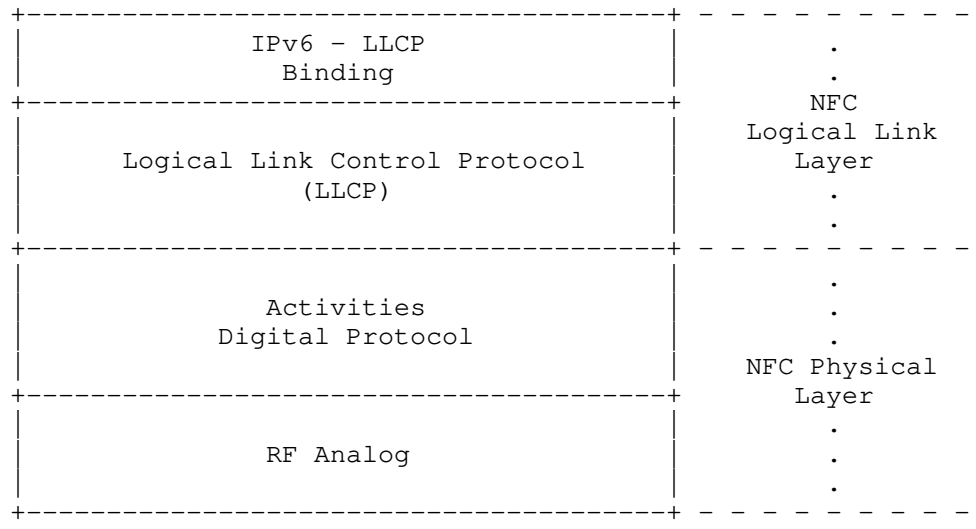

Figure 1: Protocol Stack of NFC

The LLCP consists of Logical Link Control (LLC) and MAC Mapping. The MAC Mapping integrates an existing RF protocol into the LLCP architecture. The LLC contains three components, such as Link Management, Connection-oriented Transmission, and Connectionless Transmission. The Link Management component is responsible for serializing all connection-oriented and connectionless LLC PDU (Protocol Data Unit) exchanges and for aggregation and disaggregation of small PDUs. The Connection-oriented Transmission component is responsible for maintaining all connection-oriented data exchanges

including connection set-up and termination. The Connectionless Transmission component is responsible for handling unacknowledged data exchanges.

In order to send an IPv6 packet over NFC, the packet MUST be passed down to the LLCP layer of NFC and carried by an Information Field in an LLCP Protocol Data Unit (I PDU). The LLCP does not support fragmentation and reassembly. For IPv6 addressing or address configuration, the LLCP MUST provide related information, such as link layer addresses, to its upper layer. The LLCP to IPv6 protocol binding MUST transfer the Source Service Access Point (SSAP) and Destination Service Access Point (DSAP) value to the IPv6 over NFC protocol. SSAP is a Logical Link Control (LLC) address of the source NFC-enabled device with a size of 6 bits, while DSAP means an LLC address of the destination NFC-enabled device. Thus, SSAP is a source address, and DSAP is a destination address.

3.3. NFC-enabled Device Addressing

According to NFC LLCP v1.3 [LLCP-1.3], NFC-enabled devices have two types of 6-bit addresses (i.e., SSAP and DSAP) to identify service access points. Several service access points can be installed on a NFC device. However, the SSAP and DSAP can be used as identifiers for NFC link connections with the IPv6 over NFC adaptation layer. Therefore, the SSAP can be used to generate an IPv6 interface identifier. Address values between 00h and 0Fh of SSAP and DSAP are reserved for identifying the well-known service access points, which are defined in the NFC Forum Assigned Numbers Register. Address values between 10h and 1Fh are assigned by the local LLC to services registered by local service environment. In addition, address values between 20h and 3Fh are assigned by the local LLC as a result of an upper layer service request. Therefore, the address values between 20h and 3Fh can be used for generating IPv6 interface identifiers.

3.4. MTU of NFC Link Layer

As mentioned in Section 3.2, when an IPv6 packet is transmitted, the packet MUST be passed down to LLCP of NFC and transported to an I PDU of LLCP of the NFC-enabled peer device.

The information field of an I PDU contains a single service data unit. The maximum number of octets in the information field is determined by the Maximum Information Unit (MIU) for the data link connection. The default value of the MIU for I PDUs is 128 octets. The local and remote LLCs each establish and maintain distinct MIU values for each data link connection endpoint. Also, an LLC may announce a larger MIU for a data link connection by transmitting an optional Maximum Information Unit Extension (MIUX) parameter within

the information field. If no MIUX parameter is transmitted, the MIU value is 128 bytes. Otherwise, the MTU size in NFC LLCP MUST be calculated from the MIU value as follows:

MTU = MIU = 128 + MIUX.

According to [LLCP-1.3], Figure 2 shows an example of the MIUX parameter TLV. The Type and Length fields of the MIUX parameter TLV have each a size of 1 byte. The size of the TLV Value field is 2 bytes.

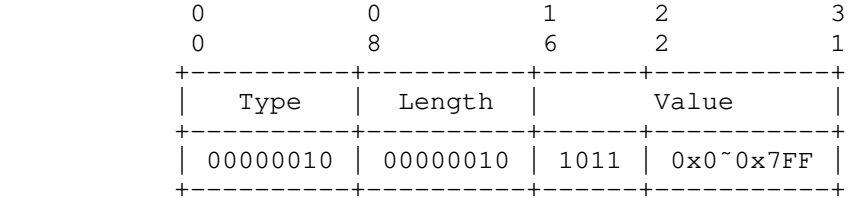

Figure 2: Example of MIUX Parameter TLV

When the MIUX parameter is used, the TLV Type field MUST be 0x02 and the TLV Length field MUST be 0x02. The MIUX parameter MUST be encoded into the least significant 11 bits of the TLV Value field. The unused bits in the TLV Value field MUST be set to zero by the sender and ignored by the receiver. The maximum possible value of the TLV Value field is 0x7FF, and the maximum size of the LLCP MTU is 2175 bytes. The MIUX value MUST be 0x480 to support the IPv6 MTU requirement (of 1280 bytes).

4. Specification of IPv6 over NFC

NFC technology has requirements owing to low power consumption and allowed protocol overhead. 6LoWPAN standards [RFC4944], [RFC6775], and [RFC6282] provide useful functionality for reducing the overhead of IPv6 over NFC. This functionality consists of link-local IPv6 addresses and stateless IPv6 address auto-configuration (see Section 4.2 and Section 4.3), Neighbor Discovery (see Section 4.4) and header compression (see Section 4.6).

4.1. Protocol Stack

Figure 3 illustrates the IPv6 over NFC protocol stack. Upper layer protocols can be transport layer protocols (e.g., TCP and UDP), application layer protocols, and others capable of running on top of IPv6.

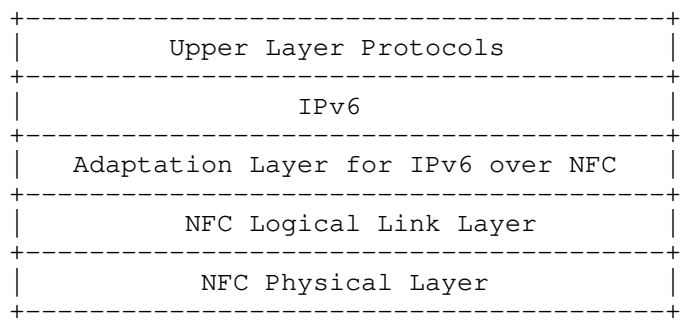

Figure 3: Protocol Stack for IPv6 over NFC

The adaptation layer for IPv6 over NFC supports neighbor discovery, stateless address auto-configuration, header compression, and fragmentation & reassembly, based on 6LoWPAN.

4.2. Stateless Address Autoconfiguration

An NFC-enabled device performs stateless address autoconfiguration as per [RFC4862]. A 64-bit Interface identifier (IID) for an NFC interface is formed by utilizing the 6-bit NFC SSAP (see Section 3.3). In the viewpoint of address configuration, such an IID should guarantee a stable IPv6 address during the course of a single connection, because each data link connection is uniquely identified by the pair of DSAP and SSAP included in the header of each LLC PDU in NFC.

Following the guidance of [RFC7136], interface identifiers of all unicast addresses for NFC-enabled devices are 64 bits long and constructed by using the generation algorithm of random (but stable) identifier (RID) [RFC7217] (see Figure 4).

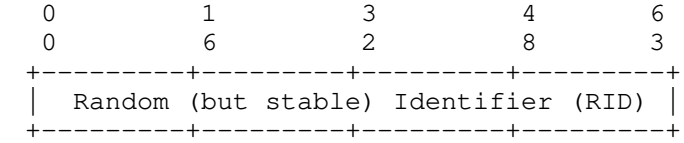

Figure 4: IID from NFC-enabled device

The RID is an output which is created by the F() algorithm with input parameters. One of the parameters is Net_Iface, and NFC Link Layer address (i.e., SSAP) is a source of the Net_Iface parameter. The 6-bit address of SSAP of NFC is short and easy to be targeted by attacks of third party (e.g., address scanning). The F() algorithm can provide secured and stable IIDs for NFC-enabled devices. In

addition, an optional parameter, Network_ID is used to increase the randomness of the generated IID.

4.3. IPv6 Link-Local Address

The IPv6 link-local address for an NFC-enabled device is formed by appending the IID to the prefix FE80::/64, as depicted in Figure 5.

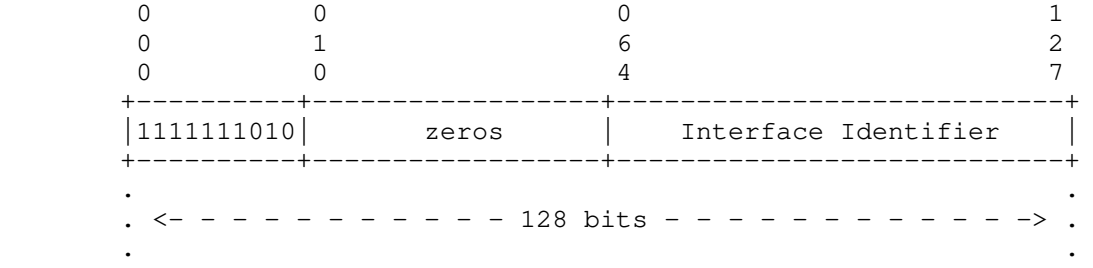

Figure 5: IPv6 link-local address in NFC

A 6LBR may obtain an IPv6 prefix for numbering the NFC network via DHCPv6 Prefix Delegation ([RFC3633]). The "Interface Identifier" can be a secured and stable IID.

4.4. Neighbor Discovery

Neighbor Discovery Optimization for 6LoWPANs ([RFC6775]) describes the neighbor discovery approach in several 6LoWPAN topologies, such as mesh topology. NFC supports mesh topologies but most of all applications would use a simple multi-hop network topology or directly connected peer-to-peer network because NFC RF range is very short.

- When an NFC-enabled 6LN is directly connected to an NFC-enabled 6LBR, the NFC 6LN MUST register its address with the 6LBR by sending a Neighbor Solicitation (NS) message with the Extended Address Registration Option (EARO) [RFC8505], and process the Neighbor Advertisement (NA) accordingly. In addition, when the 6LN and 6LBR are directly connected, DHCPv6 is used for address assignment. Therefore, Duplicate Address Detection (DAD) is not necessary between them.

- When two or more NFC devices are connected, there are two cases. One is that three or more NFC devices are linked with multi-hop connections, and the other is that they meet within a single hop range. Two NFC devices might still talk to each other (point-to-point topology), but one of them may be connected to the Internet. In a case of multi-hop topology, devices which have two or more

connections with neighbor devices, may act as routers. In a case that they meet within a single hop and they have the same properties, any of them can be a router.

- For sending Router Solicitations and processing Router Advertisements, the NFC 6LNs MUST follow Sections 5.3 and 5.4 of [RFC6775].

- When a NFC device is a 6LR or a 6LBR, the NFC device MUST follow Section 6 and 7 of [RFC6775].

4.5. Dispatch Header

All IPv6-over-NFC encapsulated datagrams are prefixed by an encapsulation header stack consisting of a Dispatch value. The only sequence currently defined for IPv6-over-NFC is the LOWPAN_IPHC compressed IPv6 header (see Section 4.6) header followed by payload, as depicted in Figure 6.

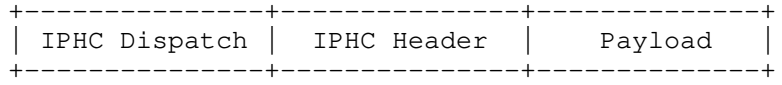

Figure 6: A IPv6-over-NFC Encapsulated 6LOWPAN_IPHC Compressed IPv6 Datagram

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

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

Figure 7: Dispatch Values

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

4.6. Header Compression

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

Therefore, IPv6 header compression in [RFC6282] MUST be implemented. Further, implementations MUST also support Generic Header Compression (GHC) of [RFC7400].

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

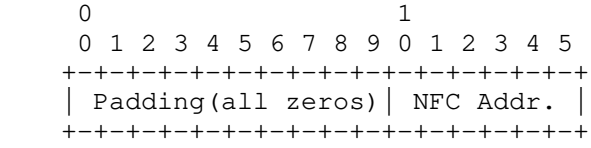

Figure 8: NFC short address format

4.7. Fragmentation and Reassembly Considerations

IIPv6-over-NFC MUST NOT use fragmentation and reassembly (FAR) at the adaptation layer for the payloads as discussed in Section 3.4. The NFC link connection for IPv6 over NFC MUST be configured with an equivalent MIU size to support the IPv6 MTU requirement (of 1280 bytes). To this end, the MIUX value is 0x480.

4.8. Unicast and Multicast Address Mapping

The address resolution procedure for mapping IPv6 non-multicast addresses into NFC link-layer addresses follows the general description in Section 4.6.1 and 7.2 of [RFC4861], unless otherwise specified.

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

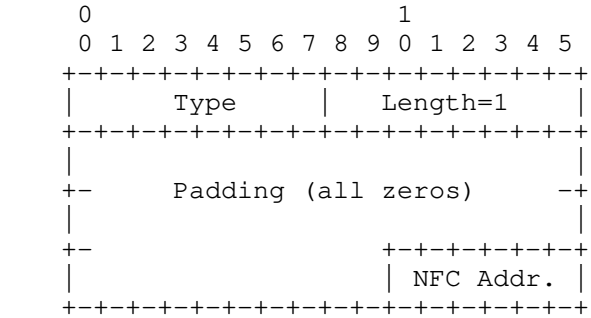

Figure 9: Unicast address mapping

Option fields:

Type:

1: for Source Link-layer address.

2: for Target Link-layer address.

Length:

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

NFC address:

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

The NFC Link Layer does not support multicast. Therefore, packets are always transmitted by unicast between two NFC-enabled devices. Even in the case where a 6LBR is attached to multiple 6LNs, the 6LBR cannot do a multicast to all the connected 6LNs. If the 6LBR needs to send a multicast packet to all its 6LNs, it has to replicate the packet and unicast it on each link.

5. Internet Connectivity Scenarios

NFC networks can either be isolated or connected to the Internet. The NFC link between two communicating devices is considered to be a point-to-point link only. An NFC link does not support a star topology or mesh network topology but only direct connections between two devices. The NFC link layer does not support packet forwarding at link layer.

5.1. NFC-enabled Device Network Connected to the Internet

Figure 10 illustrates an example of an NFC-enabled device network connected to the Internet. The distance between 6LN and 6LBR is typically 10 cm or less. For example, a laptop computer that is connected to the Internet (e.g. via Wi-Fi, Ethernet, etc.) may also support NFC and act as a 6LBR. Another NFC-enabled device may run as a 6LN and communicate with the 6LBR, as long as both are within each other's range.

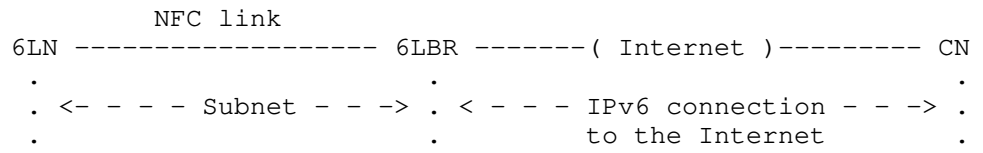

Figure 10: NFC-enabled device network connected to the Internet

Two or more 6LNs may be connected with a 6LBR, but each connection uses a different subnet. The 6LBR is acting as a router and forwarding packets between 6LNs and the Internet. Also, the 6LBR MUST ensure address collisions do not occur and forwards packets sent by one 6LN to another.

5.2. Isolated NFC-enabled Device Network

In some scenarios, the NFC-enabled device network may permanently be a simple isolated network as shown in the Figure 11.

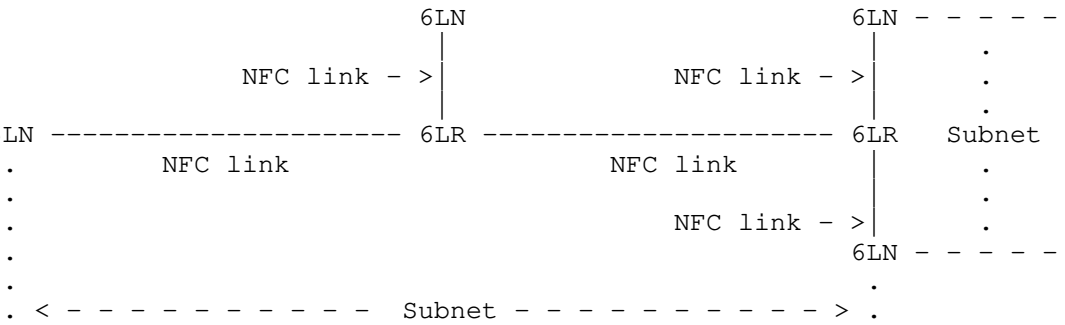

Figure 11: Isolated NFC-enabled device network

6. IANA Considerations

There are no IANA considerations related to this document.

7. Security Considerations

NFC is often considered to offer intrinsic security properties due to its short link range. When interface identifiers (IIDs) are generated, devices and users are required to consider mitigating various threats, such as correlation of activities over time, location tracking, device-specific vulnerability exploitation, and address scanning.

IPv6-over-NFC uses an IPv6 interface identifier formed from a "short address" and a set of well-known constant bits for the modified EUI-64 format. However, NFC applications use short-lived connections, and a different address is used for each connection, where the latter is of extremely short duration.

8. Acknowledgements

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

Michael Richardson, Suresh Krishnan, Pascal Thubert, Carsten Bormann, Alexandru Petrescu, James Woodyatt, Dave Thaler, Samita Chakrabarti, Gabriel Montenegro and Carles Gomez Montenegro have provided valuable feedback for this document.

9. Normative References

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

[LLCP-1.3] "NFC Logical Link Control Protocol version 1.3", NFC Forum Technical Specification , March 2016.

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

[RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, DOI 10.17487/RFC3633, December 2003, <https://www.rfc-editor.org/info/rfc3633>.

[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>.

[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>.

[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>.

[RFC7136] Carpenter, B. and S. Jiang, "Significance of IPv6 Interface Identifiers", RFC 7136, DOI 10.17487/RFC7136, February 2014, <https://www.rfc-editor.org/info/rfc7136>.

[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>.

[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>.

[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>.

Authors' Addresses

Younghwan Choi (editor)
Electronics and Telecommunications Research Institute
218 Gajeongno, Yuseung-gu
Daejeon 34129
Korea

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

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

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

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

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

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

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

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

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

6lo
Internet-Draft
Updates: 6775 (if approved)
Intended status: Standards Track
Expires: December 21, 2018

P. Thubert, Ed.
Cisco
E. Nordmark
Zededa
S. Chakrabarti
Verizon
C. Perkins
Futurewei

June 19, 2018

Registration Extensions for 6LoWPAN Neighbor Discovery
draft-ietf-6lo-rfc6775-update-21

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 Routing Registrars performing routing for host routes and/or 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 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 December 21, 2018.

Table of Contents

1. Introduction

IPv6 Low-Power Lossy Networks (LLNs) support star and mesh topologies. For such networks, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks" (6LoWPAN ND) [RFC6775] defines a registration mechanism and a central IPv6 ND Registrar to assure unique addresses. The 6LoWPAN ND mechanism reduces the dependency of the IPv6 Neighbor Discovery Protocol (IPv6 ND) [RFC4861][RFC4862] on network-layer multicast and link-layer broadcast operations.

This specification updates 6LoWPAN ND to simplify and generalizes registration in 6LoWPAN routers (6LRs). In particular, this specification modifies and extends the behavior and protocol elements of 6LoWPAN ND to enable the following actions:

- Determine the most recent location in case of node mobility
- Simplify the registration flow for Link-Local Addresses
- Support a routing-unaware Leaf Node in a Route-Over network
- Proxy registration in a Route-Over network
- Enable verification for the registration, using the Registration Ownership Verifier (ROVR)
- Registration to an IPv6 ND proxy (e.g., a Routing Registrar)
- Better support for privacy and temporary addresses

These features satisfy requirements as listed in Appendix B.

2. Terminology

2.1. BCP 14

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

2.2. References

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

- "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],
- "Problem Statement and Requirements for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing" [RFC6606], and
- "Neighbor Discovery Optimization for Low-power and Lossy Networks" [RFC6775],

2.3. Acronym Definitions

This document uses the following acronyms:

6BBR: 6LoWPAN Backbone Router

6LBR: 6LoWPAN Border Router

6LN: 6LoWPAN Node

6LR: 6LoWPAN Router

6CIO: Capability Indication Option

EARO: (Extended) Address Registration Option -- (E)ARO

EDAR: (Extended) Duplicate Address Request -- (E)DAR

EDAC: (Extended) Duplicate Address Confirmation -- (E)DAC

DAD: Duplicate Address Detection

DODAG: Destination-Oriented Directed Acyclic Graph

LLN: Low-Power and Lossy Network

NA: Neighbor Advertisement

NCE: Neighbor Cache Entry

ND: Neighbor Discovery

NDP: Neighbor Discovery Protocol

NS: Neighbor Solicitation

ROVR: Registration Ownership Verifier (pronounced rover)

RPL: IPv6 Routing Protocol for LLNs (pronounced ripple) [RFC6550]

RA: Router Advertisement

RS: Router Solicitation

TID: Transaction ID (a sequence counter in the EARO)

2.4. New Terms

Backbone Link: An IPv6 transit link that interconnects two or more Backbone Routers.

Binding: The association between an IP address, a MAC address, and other information about the node that owns the IP Address.

Registration: The process by which a 6LN registers an IPv6 Address with a 6LR in order to establish connectivity to the LLN.

Registered Node: The 6LN for which the registration is performed, according to the fields in the Extended ARO option.

Registering Node: The node that performs the registration; either the Registered Node or a proxy.

IPv6 ND Registrar: A node that can process a registration in either NS(EARO) or EDAR messages, and consequently respond with an NA or EDAC message containing the EARO and appropriate status for the registration.

Registered Address: An address registered for the Registered Node.

RFC6775-only: An implementation, a type of node, or a message that behaves only as specified by [RFC6775], as opposed to the behavior specified in this document.

Route-Over network: A network for which connectivity provided at the IP layer.

Routing Registrar: An IPv6 ND Registrar that also provides reachability services for the Registered Address, including Duplicate Address Detection and proxy Neighbor Advertisement.

Backbone Router (6BBR): A Routing Registrar that proxies the 6LoWPAN ND operations specified in this document to assure that multiple LLNs federated by a backbone link operate as a single IPv6 subnetwork.

updated: A 6LN, a 6LR, or a 6LBR that supports this specification, in contrast to an RFC6775-only device.

3. Applicability of Address Registration Options

The Address Registration Option (ARO) in [RFC6775] facilitates Duplicate Address Detection (DAD) for hosts and populates Neighbor Cache Entries (NCEs) [RFC4861] in the routers. This reduces the reliance on multicast operations, which are often as intrusive as broadcast, in IPv6 ND operations (see [I-D.ietf-mboned-ieee802-mcast-problems]).

This document specifies new status codes for registrations rejected by a 6LR or a 6LBR for reasons other than address duplication. Examples include:

- the router running out of space;
- a registration bearing a stale sequence number which could happen if the host moves after the registration was placed;
- a host misbehaving and attempting to register an invalid address such as the unspecified address [RFC4291];
- a host using an address that 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 router, depending on the returned error.

The ability to return errors to address registrations is not intended to be used to restrict the ability of hosts to form and use multiple addresses. Each host may form and register a number of addresses for enhanced privacy, using mechanisms such as "Privacy Extensions for Stateless Address Autoconfiguration (SLAAC) in IPv6" [RFC4941], and SHOULD conform to "Host Address Availability Recommendations" [RFC7934].

In IPv6 ND [RFC4861], a router needs enough storage to hold NCEs for all directly connected addresses to which it is currently forwarding packets (unused entries may be flushed). In contrast, a router serving the Address Registration mechanism needs enough storage to hold NCEs for all the addresses that may be registered to it, regardless of whether or not they are actively communicating. The number of registrations supported by a 6LoWPAN Router (6LR) or 6LoWPAN Border Router (6LBR) MUST be clearly documented by the vendor and the dynamic use of associated resources SHOULD be made available to the network operator, e.g., to a management console. Network administrators need to ensure that 6LR/6LBRs in their network support the number and type of devices that can register to them, based on the number of IPv6 addresses that those devices require and their address renewal rate and behavior.

4. Extended Neighbor Discovery Options and Messages

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

4.1. Extended Address Registration Option (EARO)

The Address Registration Option (ARO) is defined in section 4.1 of [RFC6775].

This specification introduces the Extended Address Registration Option (EARO) based on the ARO for use in NS and NA messages. The EARO includes a sequence counter called Transaction ID (TID) that is used to determine the latest location of a registering mobile device. A new 'T' flag indicates the presence of the TID field is populated and that the option is an EARO. A 6LN requests routing or proxy services from a 6LR using a new 'R' flag in the EARO.

The EUI-64 field is redefined and renamed ROVR in order to carry different types of information, e.g., cryptographic information of variable size. A larger ROVR size MAY be used if and only if backward compatibility is not an issue in the particular LLN. The length of the ROVR field expressed in units of 8 bytes is the Length of the option minus 1.

Section 5.1 discusses those changes in depth.

The format of the EARO is shown in Figure 1:

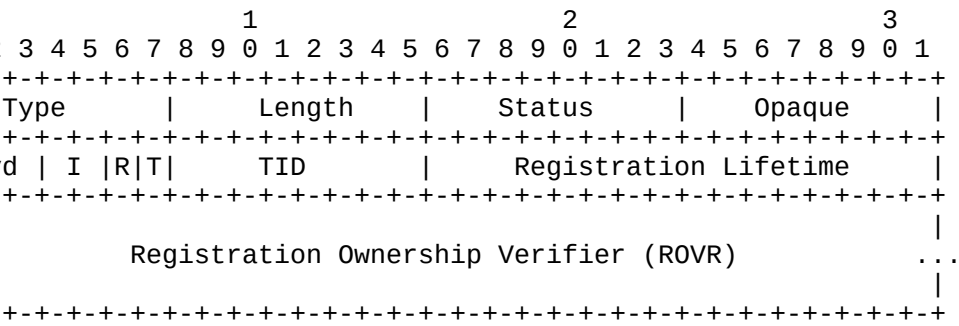

Figure 1: EARO Option Format

Option Fields:

Type: 33

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

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

Opaque: An octet opaque to ND; the 6LN MAY pass it transparently to another process. It MUST be set to zero when not used.

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

I: Two-bit Integer: A value of zero indicates that the Opaque field carries an abstract index that is used to decide in which routing topology the address is expected to be injected. In that case, the Opaque field is passed to a routing process with the indication that it carries topology information, and the value of 0 indicates default. All other values of "I" are reserved and MUST NOT be used.

R: The Registering Node sets the 'R' flag to request reachability services for the registered address from a Routing Registrar.

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

TID: One-byte unsigned integer; a Transaction ID that is maintained by the node and incremented with each transaction of one or more registrations performed at the same time to one or more 6LRs. This field MUST be ignored if the 'T' flag is not set.

Registration Lifetime: 16-bit integer; expressed in minutes. A value of 0 indicates that the registration has ended and that the associated state MUST be removed.

Registration Ownership Verifier (ROVR): Enables the correlation between multiple attempts to register a same IPv6 Address. The ROVR size MUST be 64 bits when backward compatibility is needed; otherwise the size MAY be 128, 192, or 256 bits.

Value	Description
0..2	As defined in [RFC6775]. Note: a Status of 1 ("Duplicate Address") applies to the Registered Address. If the Source Address conflicts with an existing registration, "Duplicate Source Address" MUST be used.
3	Moved: The registration failed because it is not the most recent. This Status indicates that the registration is rejected because another more recent registration was done, as indicated by a same ROVR 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 ROVR collision.
4	Removed: The binding state was removed. This status MAY be placed in an NA(EARO) message that is sent as the rejection of a proxy registration to an IPv6 ND Registrar, or in an asynchronous NA(EARO) at any time.
5	Validation Requested: The Registering Node is challenged for owning the Registered Address or for being an acceptable proxy for the registration. An IPv6 ND Registrar MAY place this Status in asynchronous DAC or NA messages.
6	Duplicate Source Address: The address used as source of the NS(EARO) conflicts with an existing registration.
7	Invalid Source Address: The address used as source of the NS(EARO) is not a Link-Local Address.
8	Registered Address topologically incorrect: The address being registered is not usable on this link.
9	6LBR Registry saturated: A new registration cannot be accepted because the 6LBR Registry is saturated. Note: this code is used by 6LBRs instead of Status 2 when responding to a Duplicate Address message exchange and is passed on to the Registering Node by the 6LR.
10	Validation Failed: The proof of ownership of the registered address is not correct.

Table 1: EARO Status

4.2. Extended Duplicate Address Message Formats

The DAR and DAC messages share a common base format as defined in section 4.4 of [RFC6775]. Those messages enable information from the ARO to be transported over multiple hops. The DAR and DAC are extended as shown in Figure 2:

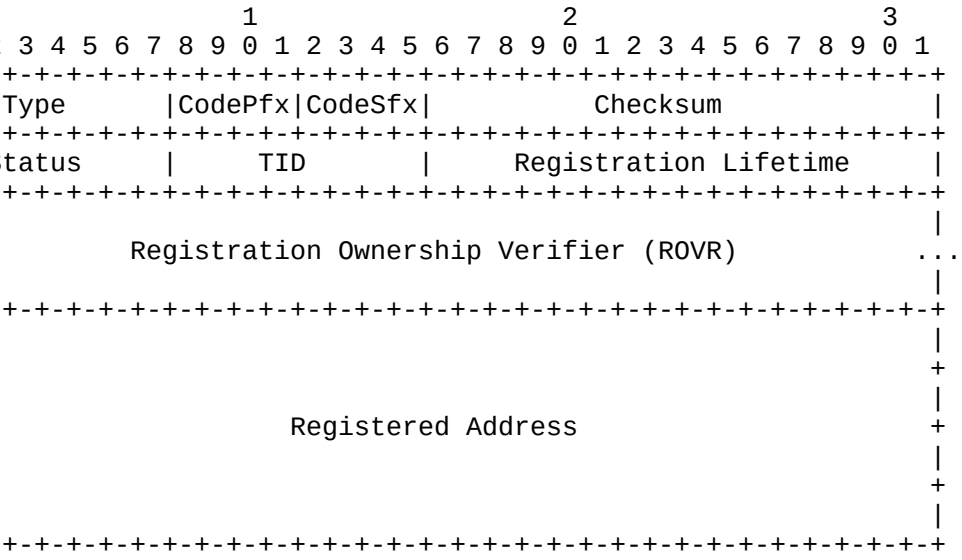

Figure 2: Duplicate Address Messages Format

Modified Message Fields:

Code: The ICMP Code [RFC4443] for Duplicate Address Messages is split in two 4-bit fields, the Code Prefix and the Code Suffix. The Code Prefix MUST be set to zero by the sender and MUST be ignored by the receiver. A non-null value of the Code Suffix indicates support for this specification. It MUST be set to 1 when operating in a backward-compatible mode, indicating a ROVR size of 64 bits. It MAY be 2, 3 or 4, denoting a ROVR size of 128, 192, and 256 bits, respectively.

TID: 1-byte integer; same definition and processing as the TID in the EARO as defined in Section 4.1. This field MUST be ignored if the ICMP Code is null.

Registration Ownership Verifier (ROVR): The size of the ROVR is known from the ICMP Code Suffix. This field has the same definition and processing as the ROVR in the EARO option as defined in Section 4.1.

4.3. Extensions to the Capability Indication Option

This specification defines 5 new capability bits for use in the 6CIO, defined by [RFC7400] for use in IPv6 ND messages.

The "E" flag indicates that EARO can be used in a registration. A 6LR that supports this specification MUST set the "E" flag.

The "D" flag indicates that the 6LBR supports EDAR and EDAC messages. A 6LR that learns the "D" flag from advertisements can then exchange EDAR and EDAC messages with the 6LBR, and it also sets the "D" flag as well as the "L" flag in the 6CIO in its own advertisements. In this way, 6LNs will be able to prefer registration with a 6LR that can make use of new 6LBR features.

The new "L", "B", and "P" flags, indicate whether a router is capable of acting as 6LR, 6LBR, and Routing Registrar (e.g., 6BBR), respectively. These flags are not mutually exclusive; an updated node can advertise multiple collocated functions.

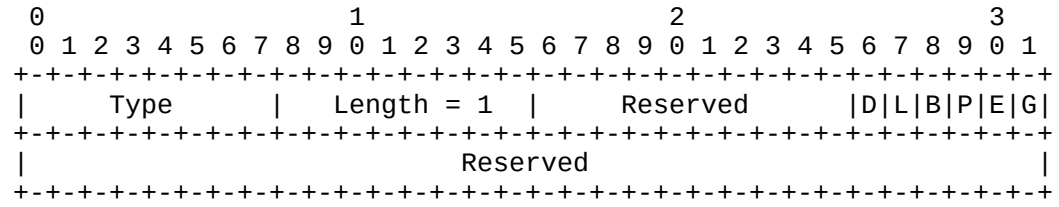

Figure 3: New Capability Bits in the 6CIO

Option Fields:

Type: 36

L: Node is a 6LR.

B: Node is a 6LBR.

P: Node is a Routing Registrar.

E: Node is an IPv6 ND Registrar -- i.e., it supports registrations based on EARO.

D: 6LBR supports EDAR and EDAC messages.

5. Updating RFC 6775

The Extended Address Registration Option (EARO) (see Section 4.1) updates the ARO used within NS and NA messages between a 6LN and a 6LR. The update enables a registration to a Routing Registrar in order to obtain additional services, such as return routability to the Registered Address by such means as routing and/or proxy Neighbor Discovery, as illustrated in Figure 4.

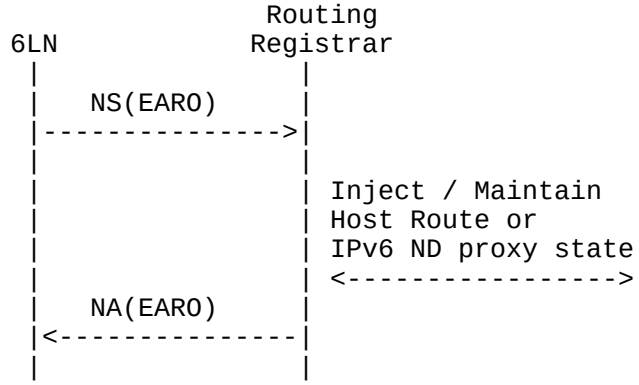

Figure 4: (Re-)Registration Flow

Similarly, EDAR and EDAC update the DAR and DAC messages so as to transport the new information between 6LRs and 6LBRs across an LLN mesh. The extensions to the ARO option are the Duplicate Address Request (DAR) and Duplicate Address Confirmation (DAC), used in the Duplicate Address messages. They convey the additional information all the way to the 6LBR.

In turn the 6LBR may proxy the registration to obtain reachability services from a Routing Registrar such as a 6BBR, as illustrated in Figure 5. This specification avoids the Duplicate Address message flow for Link-Local Addresses in a Route-Over [RFC6606] topology (see Section 5.6).

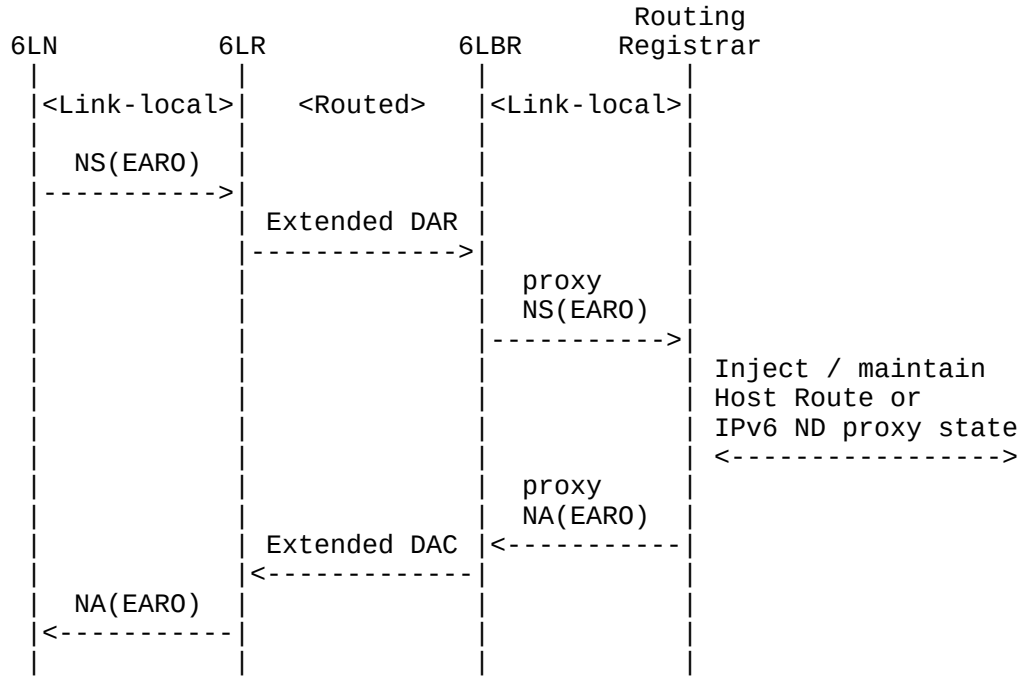

Figure 5: (Re-)Registration Flow

This specification allows multiple registrations, including for privacy / temporary addresses and provides a mechanism to help clean up stale registration state as soon as possible, e.g., after a movement (see Section 7).

Section 5 of [RFC6775] specifies how a 6LN bootstraps an interface and locates available 6LRs. A Registering Node SHOULD register to a 6LR that supports this specification if one is found, as discussed in Section 6.1, instead of registering to an RFC6775-only one; otherwise the Registering Node operates in a backward-compatible fashion when attaching to an RFC6775-only 6LR.

5.1. Extending the Address Registration Option

The Extended ARO (EARO) updates the ARO and is backward compatible with the ARO if and only if the Length of the option is set to 2. Its format is presented in Section 4.1. More details on backward compatibility can be found in Section 6.

The Neighbor Solicitation (NS) and the ARO are modified as follows:

- The Target Address in the NS containing the EARO is now the field that indicates the address that is being registered, as opposed to the Source Address field as specified in [RFC6775] (see Section 5.5). This change enables a 6LBR to send a proxy registration for a 6LN's address to a Routing Registrar, and also avoids in most cases the use of an address as source address before it is registered.

- The EUI-64 field in the ARO Option is renamed Registration Ownership Verifier (ROVR) and is not required to be derived from a MAC address (see Section 5.3).

- The option Length MAY be different than 2 and take a value between 3 and 5, in which case the EARO is not backward compatible with an ARO. The increase of size corresponds to a larger ROVR field, so the size of the ROVR is inferred from the option Length.

- A new Opaque field is introduced to carry opaque information in case the registration is relayed to another process, e.g., to be advertised by a routing protocol. A new "I" field provides a type for the opaque information, and indicates the other process to which the 6LN passes the opaque value. A value of Zero for I indicates topological information to be passed to a routing process if the registration is redistributed. In that case, a value of Zero for the Opaque field is backward-compatible with the reserved fields that are overloaded, and the meaning is to use the default topology.

- This document specifies a new flag in the EARO, the 'R' flag. If the 'R' flag is set, the Registering Node requests the 6LR to ensure reachability for the Registered Address, e.g., by means of routing or proxying ND. Conversely, when it is not set, the 'R' flag indicates that the Registering Node is a router, and that it will advertise reachability to the Registered Address via a routing protocol (such as RPL [RFC6550]).

- A node that supports this specification MUST be provide a Transaction ID (TID) field in the EARO, and set the 'T' flag to indicate the presence of the TID (see Section 5.2).

- Finally, this specification introduces new status codes to help diagnose the cause of a registration failure (see Table 1).

A 6LN that acts only as a host, when registering, MUST set the 'R' flag to indicate that it is not a router and that it will not handle its own reachability. A 6LR that manages its reachability SHOULD NOT set the 'R' flag; if it does, routes towards this router may be installed on its behalf and may interfere with those it advertises.

5.2. Transaction ID

The TID is a sequence number that is incremented by the 6LN with each re-registration to a 6LR. The TID is used to determine the recency of the registration request. The network uses the most recent TID to determine the most recent known location(s) of a moving 6LN. When a Registered Node is registered with multiple 6LRs in parallel, the same TID MUST be used. This enables the 6LBRs and/or Routing Registrars to determine whether the registrations are identical, and to distinguish that situation from a movement (for example, see Appendix A and Section 5.7).

5.2.1. Comparing TID values

The operation of the TID is fully compatible with that of the RPL Path Sequence counter as described in the "Sequence Counter Operation" section of the "IPv6 Routing Protocol for Low-Power and Lossy Networks" [RFC6550] specification.

A TID is deemed to be more recent than another when its value is greater as determined by the operations detailed in this section.

The TID range is subdivided in a 'lollipop' fashion ([Perlman83]), where the values from 128 and greater are used as a linear sequence to indicate a restart and bootstrap the counter, and the values less than or equal to 127 used as a circular sequence number space of size 128 as in [RFC1982]. Consideration is given to the mode of operation when transitioning from the linear region to the circular region. Finally, when operating in the circular region, if sequence numbers are determined to be too far apart then they are not comparable, as detailed below.

A window of comparison, SEQUENCE_WINDOW = 16, is configured based on a value of 2^N, where N is defined to be 4 in this specification.

For a given sequence counter,

1. The sequence counter SHOULD be initialized to an implementation defined value which is 128 or greater prior to use. A recommended value is 240 (256 - SEQUENCE_WINDOW).

2. When a sequence counter increment would cause the sequence counter to increment beyond its maximum value, the sequence counter MUST wrap back to zero. When incrementing a sequence counter greater than or equal to 128, the maximum value is 255. When incrementing a sequence counter less than 128, the maximum value is 127.

3. When comparing two sequence counters, the following rules MUST be applied:

    1. When a first sequence counter A is in the interval [128..255] and a second sequence counter B is in [0..127]:

        1. If (256 + B - A) is less than or equal to SEQUENCE_WINDOW, then B is greater than A, A is less than B, and the two are not equal.

        2. If (256 + B - A) is greater than SEQUENCE_WINDOW, then A is greater than B, B is less than A, and the two are not equal.

        For example, if A is 240, and B is 5, then (256 + 5 - 240) is 21. 21 is greater than SEQUENCE_WINDOW (16), thus 240 is greater than 5. As another example, if A is 250 and B is 5, then (256 + 5 - 250) is 11. 11 is less than SEQUENCE_WINDOW (16), thus 250 is less than 5.

    2. In the case where both sequence counters to be compared are less than or equal to 127, and in the case where both sequence counters to be compared are greater than or equal to 128:

        1. If the absolute magnitude of difference between the two sequence counters is less than or equal to SEQUENCE_WINDOW, then a comparison as described in [RFC1982] is used to determine the relationships greater than, less than, and equal.

        2. If the absolute magnitude of difference of the two sequence counters is greater than SEQUENCE_WINDOW, then a desynchronization has occurred and the two sequence numbers are not comparable.

4. If two sequence numbers are determined to be not comparable, i.e., the results of the comparison are not defined, then a node should give precedence to the sequence number that was most recently incremented. Failing this, the node should select the sequence number in order to minimize the resulting changes to its own state.

5.3. Registration Ownership Verifier (ROVR)

The ROVR field replaces the EUI-64 field of the ARO defined in [RFC6775]. It is associated in the 6LR and the 6LBR with the registration state. The ROVR can be a unique ID of the Registering

Node, such as the EUI-64 address of an interface. This can also be a token obtained with cryptographic methods which can be used in additional protocol exchanges to associate a cryptographic identity (key) with this registration to ensure that only the owner can modify it later, if the proof-of-ownership of the ROVR can be obtained (more in Section 5.6). The scope of a ROVR is the registration of a particular IPv6 Address and it MUST NOT be used to correlate registrations of different addresses.

The ROVR can be of different types; the type is signaled in the message that carries the new type. For instance, the type can be a cryptographic string and used to prove the ownership of the registration as specified in "Address Protected Neighbor Discovery for Low-power and Lossy Networks" [I-D.ietf-6lo-ap-nd]. In order to support the flows related to the proof-of-ownership, this specification introduces new status codes "Validation Requested" and "Validation Failed" in the EARO.

Note on ROVR collision: different techniques for forming the ROVR will operate in different name-spaces. [RFC6775] operates on EUI-64(TM) addresses. [I-D.ietf-6lo-ap-nd] generates cryptographic tokens. While collisions are not expected in the EUI-64 name-space only, they may happen in the case of [I-D.ietf-6lo-ap-nd] and in a mixed situation. An implementation that understands the name-space MUST consider that ROVRs from different name-spaces are different even if they have the same value. An RFC6775-only 6LR or 6LBR will confuse the name-spaces, which slightly increases the risk of a ROVR collision. A collision of ROVR has no effect if the two Registering Nodes register different addresses, since the ROVR is only significant within the context of one registration. A ROVR is not expected to be unique to one registration, as this specification allows a node to use the same ROVR to register multiple IPv6 addresses. This is why the ROVR MUST NOT be used as a key to identify the Registering Node, or as an index to the registration. It is only used as a match to ensure that the node that updates a registration for an IPv6 address is the node that made the original registration for that IPv6 address. Also, when the ROVR is not an EUI-64 address, then it MUST NOT be used as the interface ID of the Registered Address. This way, a registration that uses that ROVR will not collide with that of an IPv6 Address derived from EUI-64 and using the EUI-64 as ROVR per [RFC6775].

The Registering Node SHOULD store the ROVR, or enough information to regenerate it, in persistent memory. If this is not done and an event such as a reboot causes a loss of state, re-registering the same address could be impossible until the 6LRs and the 6LBR time out the previous registration, or a management action is taken to clear the relevant state in the network.

5.4. Extended Duplicate Address Messages

In order to map the new EARO content in the Extended Duplicate Address (EDA) messages, a new TID field is added to the Extended DAR (EDAR) and the Extended DAC (EDAC) messages as a replacement of the Reserved field, and a non-null value of the ICMP Code indicates support for this specification. The format of the EDAR and EDAC messages is presented in Section 4.2.

As with the EARO, the Extended Duplicate Address messages are backward compatible with the RFC6775-only versions as long as the ROVR field is 64 bits long. Remarks concerning backwards compatibility for the protocol between the 6LN and the 6LR apply similarly between a 6LR and a 6LBR.

5.5. Registering the Target Address

An NS message with an EARO is a registration if and only if it also carries an SLLA Option [RFC6775]. The EARO can also be used in NS and NA messages between Routing Registrars to determine the distributed registration state; in that case, it does not carry the SLLA Option and is not confused with a registration.

The Registering Node is the node that performs the registration to the Routing Registrar. As in [RFC6775], it may be the Registered Node as well, in which case it registers one of its own addresses and indicates its own MAC Address as Source Link Layer Address (SLLA) in the NS(EARO).

This specification adds the capability to proxy the registration operation on behalf of a Registered Node that is reachable over an LLN mesh. In that case, if the Registered Node is reachable from the Routing Registrar via a Mesh-Under mesh, the Registering Node indicates the MAC Address of the Registered Node as the SLLA in the NS(EARO). If the Registered Node is reachable over a Route-Over mesh from the Registering Node, the SLLA in the NS(ARO) is that of the Registering Node. This enables the Registering Node to attract the packets from the Routing Registrar and route them over the LLN to the Registered Node.

In order to enable the latter operation, 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 field. With this convention, a TLLA option indicates the link-layer address of the 6LN that owns the address.

A Registering Node (e.g., a 6LBR also acting as RPL Root) that advertises reachability for the 6LN MUST place its own Link Layer Address in the SLLA Option of the registration NS(EARO) message. This maintains compatibility with RFC6775-only 6LoWPAN ND [RFC6775].

5.6. Link-Local Addresses and Registration

LLN nodes are often not wired and may move. There is no guarantee that a Link-Local Address remain unique among a huge and potentially variable set of neighboring nodes.

Compared to [RFC6775], this specification only requires that a Link-Local Address be unique from the perspective of the two nodes that use it to communicate (e.g., the 6LN and the 6LR in an NS/NA exchange). This simplifies the DAD process in a Route-Over topology for Link-Local Addresses by avoiding an exchange of EDA messages between the 6LR and a 6LBR for those addresses.

An exchange between two nodes using Link-Local Addresses implies that they are reachable over one hop. A node MUST register a Link-Local Address to a 6LR in order to obtain further reachability by way of that 6LR, 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 a previously registered address, then the Link-Local Address is unique from the standpoint of this 6LR and the registration is not a duplicate. Two different 6LRs might claim the same Link-Local Address but different link-layer addresses. In that case, a 6LN MUST only interact with at most one of the 6LRs.

The exchange of EDAR and EDAC messages between the 6LR and a 6LBR, which ensures that an address is unique across the domain covered by the 6LBR, does not need to take place for Link-Local Addresses.

When sending an NS(EARO) to a 6LR, a 6LN 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 an address that is already registered to the 6LR, or the address that is being registered.

When a 6LN starts up, it typically multicasts a RS and receives one or more unicast RA messages from 6LRs. If the 6LR can process EARO messages, then it places a 6CIO in its RA message with the "E" Flag set as required in Section 6.1.

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 an address for which DAD is not required (see [RFC6775]), e.g., derived from a globally unique EUI-64 address; using the SLLA Option in the NS is consistent with existing ND specifications such as the "Optimistic Duplicate Address Detection (ODAD) for IPv6" [RFC4429]. The 6LN MAY then use that address to register one or more other addresses.

A 6LR that supports this specification replies with an NA(EARO), setting the appropriate status. Since there is no exchange of EDAR or EDAC messages 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 is placed in the NS(EARO) message as required in [RFC6775].

A node registers its IPv6 Global Unicast Addresses (GUAs) to a 6LR in order to establish global reachability for these addresses via that 6LR. When registering with an updated 6LR, a Registering Node does not use a GUA as Source Address, in contrast to a node that complies to [RFC6775]. For non-Link-Local Addresses, the exchange of EDAR and EDAC messages MUST conform to [RFC6775], but the extended formats described in this specification for the DAR and the DAC are used to relay the extended information in the case of an EARO.

5.7. 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; as discussed in Section 5.6, this 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. 6LBRs and Routing Registrars may store additional registration information and use synchronization protocols that are out of scope of this document.

A 6LR cannot accept a new registration when its registration storage space is exhausted. In that situation, the EARO is returned in an NA message with a Status Code of "Neighbor Cache Full" (Table 1), and the Registering Node may attempt to register to another 6LR.

If the registry in the 6LBR is full, then the 6LBR cannot decide whether a registration for a new address is a duplicate. In that case, the 6LBR replies to an EDAR message with an EDAC message that carries a new Status Code indicating "6LBR Registry Saturated" (Table 1). Note: this code is used by 6LBRs instead of "Neighbor Cache Full" when responding to a Duplicate Address message exchange and is passed on to the Registering Node by the 6LR. There is no point for the node to retry this registration via another 6LR, since

the problem is network-wide. The node may either abandon that address, de-register other addresses first to make room, or keep the address in TENTATIVE state and retry later.

A node renews an existing registration by sending a new NS(EARO) message for the Registered Address, and the 6LR MUST report the new registration to the 6LBR.

A node that ceases to use an address SHOULD attempt to de-register that address from all the 6LRs to which it has registered the address. This is achieved using an NS(EARO) message with a Registration Lifetime of 0. If this is not done, the associated state will remain in the network till the current Registration Lifetime expires and this may lead to a situation where the 6LR resources become saturated, even if they are correctly planned to start with. The 6LR may then take defensive measures that may prevent this node or some other nodes from owning as many addresses as they request (see Section 7).

A node that moves away from a particular 6LR SHOULD attempt to de-register all of its addresses registered to that 6LR and register to a new 6LR with an incremented TID. When/if the node appears elsewhere, an asynchronous NA(EARO) or EDAC message with a Status Code of "Moved" SHOULD be used to clean up the state in the previous location. The "Moved" status can be used by a Routing Registrar in an NA(EARO) message to indicate that the ownership of the proxy state was transferred to another Routing Registrar due to movement of the device. If the receiver of the message has registration state corresponding to the related address, it SHOULD propagate the status down the forwarding path to the Registered Node (e.g., reversing an existing RPL [RFC6550] path as prescribed in [I-D.ietf-roll-efficient-npdao]). Whether it could do so or not, the receiver MUST clean up said state.

Upon receiving an NS(EARO) message with a Registration Lifetime of 0 and determining that this EARO is the most recent for a given NCE (see Section 5.2), a 6LR cleans up its NCE. If the address was registered to the 6LBR, then the 6LR MUST report to the 6LBR, through a Duplicate Address exchange with the 6LBR, indicating the null Registration Lifetime and the latest TID that this 6LR is aware of.

Upon receiving the EDAR message, the 6LBR evaluates if this is the most recent TID it has received for that particular registry entry. If so, then the EDAR is answered with an EDAC message bearing a Status of "Success" and the entry is scheduled to be removed. Otherwise, a Status Code of "Moved" is returned instead, and the existing entry is maintained.

When an address is scheduled to be removed, the 6LBR SHOULD keep its NCE in a DELAY state [RFC4861] for a configurable period of time, so as to protect a mobile node that de-registered from one 6LR and did not register yet to a new one, or the new registration did not yet reach the 6LBR due to propagation delays in the network. Once the DELAY time is passed, the 6LBR silently removes its entry.

6. Backward Compatibility

This specification changes the behavior of the peers in a registration flow. To enable backward compatibility, a 6LN that registers to a 6LR that is not known to support this specification MUST behave in a manner that is backward-compatible with [RFC6775]. On the contrary, if the 6LR is found to support this specification, then the 6LN MUST conform to this specification when communicating with that 6LR.

A 6LN that supports this specification MUST always use an EARO as a replacement for an ARO in its registration to a router. This is backward-compatible since the 'T' flag and TID field are reserved in [RFC6775], and are ignored by an RFC6775-only router. A router that supports this specification MUST answer an NS(ARO) and an NS(EARO) with an NA(EARO). A router that does not support this specification will consider the ROVR as an EUI-64 address and treat it the same, which has no consequence if the Registered Addresses are different.

6.1. Signaling EARO Support

"Generic Header Compression for IPv6 over 6LoWPANs" [RFC7400] specifies the 6LoWPAN Capability Indication Option (6CIO) to indicate a node's capabilities to its peers. The 6CIO MUST be present in both Router Solicitation (RS) and Router Advertisement (RA) messages, unless the 6CIO information was already shared in recent exchanges, or pre-configured in all nodes in a network. In any case, a 6CIO MUST be placed in an RA message that is sent in response to an RS with a 6CIO.

Section 4.3 defines a new flag for the 6CIO to signal support for EARO by the issuer of the message. New flags are also added to the 6CIO to signal the sender's capability to act as a 6LR, 6LBR, and Routing Registrar (see Section 4.3).

Section 4.3 also defines a new flag that indicates the support of EDAR and EDAC messages by the 6LBR. This flag is valid in RA messages but not in RS messages. More information on the 6LBR is found in a separate Authoritative Border Router Option (ABRO). The ABRO is placed in RA messages as prescribed by [RFC6775]; in particular, it MUST be placed in an RA message that is sent in

response to an RS with a 6CIO indicating the capability to act as a 6LR, since the RA propagates information between routers.

6.2. RFC6775-only 6LN

An RFC6775-only 6LN will use the Registered Address as the source address of the NS message and will not use an EARO. An updated 6LR MUST accept that registration if it is valid per [RFC6775], and it MUST manage the binding cache accordingly. The updated 6LR MUST then use the RFC6775-only DAR and DAC messages as specified in [RFC6775] to indicate to the 6LBR that the TID is not present in the messages.

The main difference from [RFC6775] is that the exchange of DAR and DAC messages for the purpose of DAD is avoided for Link-Local Addresses. In any case, the 6LR MUST use an EARO in the reply, and can use any of the Status codes defined in this specification.

6.3. RFC6775-only 6LR

An updated 6LN discovers the capabilities of the 6LR in the 6CIO in RA messages from that 6LR; if the 6CIO was not present in the RA, then the 6LR is assumed to be a RFC6775-only 6LR.

An updated 6LN MUST use an EARO in the request regardless of the type of 6LR, RFC6775-only or updated, which implies that the 'T' flag is set. It MUST use a ROVR of 64 bits if the 6LR is an RFC6775-only 6LR.

If an updated 6LN moves from an updated 6LR to an RFC6775-only 6LR, the RFC6775-only 6LR will send an RFC6775-only DAR message, which cannot be compared with an updated one for recency. Allowing RFC6775-only DAR messages to update a state established by the updated protocol in the 6LBR would be an attack vector and that cannot be the default behavior. But if RFC6775-only and updated 6LRs coexist temporarily in a network, then it makes sense for an administrator to install a policy that allows this, using some method out of scope for this document.

6.4. RFC6775-only 6LBR

With this specification, the Duplicate Address messages are extended to transport the EARO information. As with the NS/NA exchange, an updated 6LBR MUST always use the EDAR and EDAC messages.

Note that an RFC6775-only 6LBR will accept and process an EDAR message as if it were an RFC6775-only DAR, as long as the ROVR is 64 bits long. An updated 6LR discovers the capabilities of the 6LBR in

the 6CIO in RA messages from the 6LR; if the 6CIO was not present in any RA, then the 6LBR is assumed to be a RFC6775-only 6LBR.

If the 6LBR is RFC6775-only, the 6LR MUST use only the 64 leftmost bits of the ROVR, and place the result in the EDAR message to maintain compatibility. This way, the support of DAD is preserved.

7. Security Considerations

This specification extends [RFC6775], and the security section of that document also applies to this document. In particular, the link layer SHOULD be sufficiently protected to prevent rogue access.

[RFC6775] does not protect the content of its messages and expects a lower layer encryption to defeat potential attacks. This specification requires the LLN MAC to provide secure unicast to/from a Routing Registrar and secure Broadcast or Multicast from the Routing Registrar in a way that prevents tampering with or replaying the Neighbor Discovery messages.

This specification recommends using privacy techniques (see Section 8), and protecting against address theft by methods outside the scope of this document. As an example, "Address Protected Neighbor Discovery for Low-power and Lossy Networks" [I-D.ietf-6lo-ap-nd] guarantees the ownership of the Registered Address using a cryptographic ROVR.

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 registrations, which effectively denies the requesting node the capability to use a new address. In order to alleviate those concerns, Section 5.7 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:

- A node that ceases to use an address SHOULD attempt to de-register that address from all the 6LRs to which it is registered. See Section 5.2 for the mechanism to avoid replay attacks and avoiding the use of stale registration information.

- The Registration lifetimes SHOULD be individually configurable for each address or group of addresses. 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

larger yet of a same order as the duration of the expectation of presence.

- The router (6LR or 6LBR) SHOULD be configurable so as to limit the number of addresses that can be registered by a single node, but as a protective measure only. In any case, a router MUST be able to keep a minimum number of addresses per node. That minimum depends on the type of device and ranges between 3 for a very constrained LLN and 10 for a larger device. A node may be identified by its MAC address, as long as it is not obfuscated by privacy measures. A stronger identification (e.g., by security credentials) is RECOMMENDED. When the maximum is reached, the router SHOULD use a Least-Recently-Used (LRU) algorithm to clean up the addresses, keeping at least one Link-Local Address. The router SHOULD attempt to keep one or more stable addresses if stability can be determined, e.g., because they are used over a much longer time span than other (privacy, shorter-lived) addresses.

- In order to avoid denial of registration for the lack of resources, administrators should take great care to deploy adequate numbers of 6LRs 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 an LLN is a more capable node than the average 6LR, but in a network condition where it may become saturated, a particular LLN should distribute the 6LBR functionality, for instance by leveraging a high speed Backbone Link and Routing Registrars to aggregate multiple LLNs into a larger subnet.

The LLN nodes depend on a 6LBR and may use the services of a routing Registrar for their operation. A trust model MUST be put in place to ensure that only authorized 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. This trust model could be at a minimum based on a Layer-2 access control, or could provide role validation as well (see Req5.1 in Appendix B.5).

8. Privacy Considerations

As indicated in Section 3, this protocol does not limit the number of IPv6 addresses that each device can form. However, to mitigate denial-of-service attacks, it can be useful as a protective measure to have a limit that is high enough not to interfere with the normal behavior of devices in the network. 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 discourages using EUI-64 for forming the Interface ID in the Link-Local Address because this method prevents the usage of "SEcure Neighbor Discovery (SEND)" [RFC3971], "Cryptographically Generated Addresses (CGA)" [RFC3972], and other address privacy techniques.

"Privacy Considerations for IPv6 Adaptation-Layer Mechanisms" [RFC8065] explains why privacy is important and how to form privacy-aware addresses. All implementations and deployments must consider the option of privacy addresses in their own environments.

The IPv6 address of the 6LN in the IPv6 header can be compressed statelessly when the Interface Identifier in the IPv6 address can be derived from the Lower Layer address. When it is not critical to benefit from that compression, e.g., the address can be compressed statefully, or it is rarely used and/or it is used only over one hop, then privacy concerns should be considered. In particular, new implementations should follow the IETF "Recommendation on Stable IPv6 Interface Identifiers" [RFC8064]. [RFC8064] recommends the use of "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)" [RFC7217] for generating Interface Identifiers to be used in SLAAC.

9. IANA Considerations

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

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

9.1. ARO Flags

IANA is requested to create a new subregistry for "ARO Flags" under the "Internet Control Message Protocol version 6 (ICMPv6) [RFC4443] Parameters".

This specification defines 8 positions, bit 0 to bit 7, and assigns bit 6 for the ’R’ flag and bit 7 for the ’T’ flag (see Section 4.1). The policy is "IETF Review" or "IESG Approval" [RFC8126].

The initial content of the registry is as shown in Table 2.

ARO Status	Description	Document
0..5	Unassigned	
6	'R' Flag	This RFC
7	'T' Flag	This RFC

Table 2: New ARO Flags

9.2. EARO I-Field

IANA is requested to create a new subregistry for "ARO Flags" under the "Internet Control Message Protocol version 6 (ICMPv6) [RFC4443] Parameters".

This specification defines 4 integer values from 0 to 3, and assigns value 0 (see Section 4.1). The policy is "IETF Review" or "IESG Approval" [RFC8126].

The initial content of the registry is as shown in Table 3.

Value	Meaning	Reference
0	Abstract Index for Topology Selection	This RFC
1..3	Unassigned	

Table 3: New subregistry for the EARO "I" Field

9.3. ICMP Codes

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

The new subregistries relate to the ICMP type 157, Duplicate Address Request (shown in Table 4), and 158, Duplicate Address Confirmation (shown in Table 5), respectively. For those two ICMP types, the ICMP Code field is split into 2 subfields, the "Code Prefix" and the "Code Suffix". The new subregistries relate to the "Code Suffix" portion of the ICMP Code. The range of "Code Suffix" is 0..15 in all cases.

The policy is "IETF Review" or "IESG Approval" [RFC8126] for both subregistries.

The new subregistries are to be initialized as follows:

Code Suffix	Meaning	Reference
0	RFC6775 DAR message	RFC 6775
1	EDAR message with 64-bit ROVR field	This RFC
2	EDAR message with 128-bit ROVR field	This RFC
3	EDAR message with 192-bit ROVR field	This RFC
4	EDAR message with 256-bit ROVR field	This RFC
5...15	Unassigned	

Table 4: New Code Suffixes for ICMP type 157 DAR message

Code Suffix	Meaning	Reference
0	RFC6775 DAC message	RFC 6775
1	EDAC message with 64-bit ROVR field	This RFC
2	EDAC message with 128-bit ROVR field	This RFC
3	EDAC message with 192-bit ROVR field	This RFC
4	EDAC message with 256-bit ROVR field	This RFC
5...15	Unassigned	

Table 5: New Code Suffixes for ICMP type 158 DAC message

9.4. New ARO Status values

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

ARO Status	Description	Document
3	Moved	This RFC
4	Removed	This RFC
5	Validation Requested	This RFC
6	Duplicate Source Address	This RFC
7	Invalid Source Address	This RFC
8	Registered Address topologically incorrect	This RFC
9	6LBR Registry saturated	This RFC
10	Validation Failed	This RFC

Table 6: New ARO Status values

9.5. New 6LoWPAN Capability Bits

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

Capability Bit	Description	Document
10	EDA Support (D bit)	This RFC
11	6LR capable (L bit)	This RFC
12	6LBR capable (B bit)	This RFC
13	Routing Registrar (P bit)	This RFC
14	EARO support (E bit)	This RFC

Table 7: New 6LoWPAN Capability Bits

10. Acknowledgments

Kudos to Eric Levy-Abegnoli who designed the First Hop Security infrastructure upon which the first backbone router was implemented. Many thanks to Sedat Gormus, Rahul Jadhav, Tim Chown, Juergen Schoenwaelder, Chris Lonvick, Dave Thaler, Adrian Farrel, Peter Yee, Warren Kumari, Benjamin Kaduk, Mirja Kuhlewind, Ben Campbell, Eric Rescorla, and Lorenzo Colitti for their various contributions and reviews. Also, many thanks to Thomas Watteyne for the world first implementation of a 6LN that was instrumental to the early tests of the 6LR, 6LBR and Backbone Router.

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, <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>.

[RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <https://www.rfc-editor.org/info/rfc4443>.

[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>.

[RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <https://www.rfc-editor.org/info/rfc6282>.

[RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <https://www.rfc-editor.org/info/rfc6775>.

[RFC7400] Bormann, C., "6LoWPAN-GHC: Generic Header Compression for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 7400, DOI 10.17487/RFC7400, November 2014, <https://www.rfc-editor.org/info/rfc7400>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://www.rfc-editor.org/info/rfc8126>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

11.2. Terminology Related References

[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>.

[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>.

11.3. 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.hou-6lo-plc]
Hou, J., Hong, Y., and X. Tang, "Transmission of IPv6 Packets over PLC Networks", draft-hou-6lo-plc-03 (work in progress), December 2017.

[I-D.ietf-6lo-ap-nd]
Thubert, P., Sarikaya, B., and M. Sethi, "Address Protected Neighbor Discovery for Low-power and Lossy Networks", draft-ietf-6lo-ap-nd-06 (work in progress), February 2018.

[I-D.ietf-6lo-backbone-router]
Thubert, P., "IPv6 Backbone Router", draft-ietf-6lo-backbone-router-06 (work in progress), February 2018.

[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-09 (work in progress), January 2018.

[I-D.ietf-6tisch-architecture]
Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", draft-ietf-6tisch-architecture-14 (work in progress), April 2018.

[I-D.ietf-mboned-ieee802-mcast-problems]
Perkins, C., McBride, M., Stanley, D., Kumari, W., and J. Zuniga, "Multicast Considerations over IEEE 802 Wireless Media", draft-ietf-mboned-ieee802-mcast-problems-01 (work in progress), February 2018.

[I-D.ietf-roll-efficient-npdao]
Jadhav, R., Thubert, P., Sahoo, R., and Z. Cao, "Efficient Route Invalidation", draft-ietf-roll-efficient-npdao-03 (work in progress), March 2018.

[I-D.struik-lwip-curve-representations]
Struik, R., "Alternative Elliptic Curve Representations", draft-struik-lwip-curve-representations-00 (work in progress), October 2017.

[I-D.thubert-roll-unaware-leaves]
Thubert, P., "Routing for RPL Leaves", draft-thubert-roll-unaware-leaves-05 (work in progress), May 2018.

[RFC1958] Carpenter, B., Ed., "Architectural Principles of the Internet", RFC 1958, DOI 10.17487/RFC1958, June 1996, <https://www.rfc-editor.org/info/rfc1958>.

[RFC1982] Elz, R. and R. Bush, "Serial Number Arithmetic", RFC 1982, DOI 10.17487/RFC1982, August 1996, <https://www.rfc-editor.org/info/rfc1982>.

[RFC3610] Whiting, D., Housley, R., and N. Ferguson, "Counter with CBC-MAC (CCM)", RFC 3610, DOI 10.17487/RFC3610, September 2003, <https://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, <https://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, <https://www.rfc-editor.org/info/rfc3971>.

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

[RFC4429] Moore, N., "Optimistic Duplicate Address Detection (DAD) for IPv6", RFC 4429, DOI 10.17487/RFC4429, April 2006, <https://www.rfc-editor.org/info/rfc4429>.

[RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, DOI 10.17487/RFC4941, September 2007, <https://www.rfc-editor.org/info/rfc4941>.

[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>.

[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>.

[RFC7428] Brandt, A. and J. Buron, "Transmission of IPv6 Packets over ITU-T G.9959 Networks", RFC 7428, DOI 10.17487/RFC7428, February 2015, <https://www.rfc-editor.org/info/rfc7428>.

[RFC7668] Nieminen, J., Savolainen, T., Isomaki, M., Patil, B., Shelby, Z., and C. Gomez, "IPv6 over BLUETOOTH(R) Low Energy", RFC 7668, DOI 10.17487/RFC7668, October 2015, <https://www.rfc-editor.org/info/rfc7668>.

[RFC7934] Colitti, L., Cerf, V., Cheshire, S., and D. Schinazi, "Host Address Availability Recommendations", BCP 204, RFC 7934, DOI 10.17487/RFC7934, July 2016, <https://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, <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>.

[RFC8105] Mariager, P., Petersen, J., Ed., Shelby, Z., Van de Logt, M., and D. Barthel, "Transmission of IPv6 Packets over Digital Enhanced Cordless Telecommunications (DECT) Ultra Low Energy (ULE)", RFC 8105, DOI 10.17487/RFC8105, May 2017, <https://www.rfc-editor.org/info/rfc8105>.

[RFC8163] Lynn, K., Ed., Martocci, J., Neilson, C., and S. Donaldson, "Transmission of IPv6 over Master-Slave/Token-Passing (MS/TP) Networks", RFC 8163, DOI 10.17487/RFC8163, May 2017, <https://www.rfc-editor.org/info/rfc8163>.

[RFC8279] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast Using Bit Index Explicit Replication (BIER)", RFC 8279, DOI 10.17487/RFC8279, November 2017, <https://www.rfc-editor.org/info/rfc8279>.

11.4. External Informative References

[IEEEstd802154] IEEE, "IEEE Standard for Low-Rate Wireless Networks", IEEE Standard 802.15.4, DOI 10.1109/IEEE P802.15.4-REVd/D01, June 2017, <http://ieeexplore.ieee.org/document/7460875/>.

[Perlman83]
Perlman, R., "Fault-Tolerant Broadcast of Routing Information", North-Holland Computer Networks 7: 395-405, 1983, <http://www.cs.illinois.edu/~pbg/courses/cs598fa09/readings/p83.pdf>.

Appendix A. Applicability and Requirements Served (Not Normative)

This specification extends 6LoWPAN ND to provide a sequence number to the registration and serves the requirements expressed in Appendix B.1 by enabling the mobility of devices from one LLN to the next. A full specification for enabling mobility based on the use of the EARO and the registration procedures defined in this document can be found in a companion document "IPv6 Backbone Router" [I-D.ietf-6lo-backbone-router]. The 6BBR is an example of a Routing Registrar that acts as an IPv6 ND proxy over a Backbone Link that federates multiple LLNs as well as the Backbone Link intself into a single IPv6 subnet. The expected registration flow in that case is illustrated in Figure 6, noting that any combination of 6LR, 6LBR and 6BBR may be collocated.

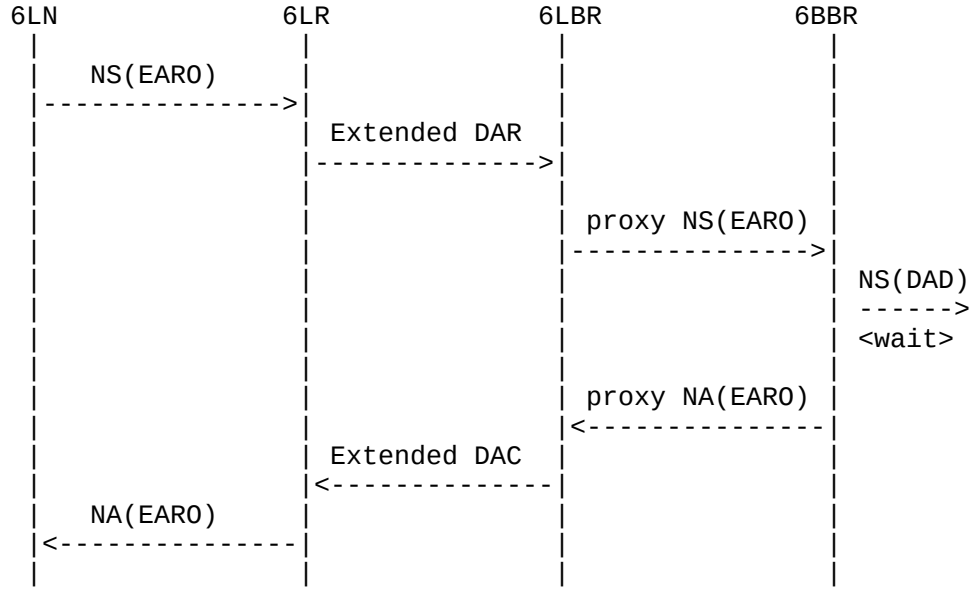

Figure 6: (Re-)Registration Flow

"6TiSCH architecture" [I-D.ietf-6tisch-architecture] describes how a 6LoWPAN ND host using the Timeslotted Channel Hopping (TSCH) mode of IEEE Std. 802.15.4 [IEEEstd802154] can connect to the Internet via a RPL mesh network. Doing so requires additions to the 6LoWPAN ND

protocol to support mobility and reachability in a secure and manageable network environment. This document specifies those new operations, and fulfills the requirements listed in Appendix B.2.

The term LLN is used loosely in this document, and intended to cover multiple types of WLANs and WPANs, including Low-Power IEEE Std. 802.11 networking, Bluetooth 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 register its IPv6 addresses with a Routing Registrar and to obtain routing services including proxy-ND operations over a Backbone Link. This satisfies the the requirements expressed in Appendix B.4.

This specification is extended by "Address Protected Neighbor Discovery for Low-power and Lossy Networks" [I-D.ietf-6lo-ap-nd] to provide a solution to some of the security-related requirements expressed in Appendix B.5.

"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 IPv6 ND ([RFC4861], [RFC4862]) and affect the operation of the wireless medium [I-D.ietf-mboned-ieee802-mcast-problems]. This serves the scalability requirements listed in Appendix B.6.

Appendix B. Requirements (Not Normative)

This section lists requirements that were discussed by the 6lo WG for an update to 6LoWPAN ND. How those requirements are matched with existing specifications at the time of this writing is shown in Appendix B.8.

B.1. Requirements Related to Mobility

Due to the unstable nature of LLN links, even in an LLN of immobile nodes, a 6LN may change its point of attachment from 6LR-a to 6LR-b, 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, e.g., by using some signaling upon the detection of the movement, or using a keep-alive mechanism with a period that is consistent with the application needs.

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

Req1.2: For that purpose, the protocol MUST enable differentiating 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 able to register its Address concurrently to multiple 6LRs.

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 an LLN can be based on RPL, which is the routing protocol that was defined by the IETF for this particular purpose. Other routing protocols are also considered by Standards Development Organizations (SDO) on the basis of the expected network characteristics. It is required that a 6LN attached via ND to a 6LR indicates whether it participates in the selected routing protocol to obtain reachability via the 6LR, or whether it expects the 6LR to manage its reachability.

The specified updates enable other specifications to define new services such as Source Address Validation (SAVI) with [I-D.ietf-6lo-ap-nd], participation as an unaware leaf to a routing protocol such as the "Routing Protocol for Low Power and Lossy Networks" [RFC6550] (RPL) with [I-D.thubert-roll-unaware-leaves], and registration to a backbone routers performing proxy Neighbor Discovery in a Low-Power and Lossy Network (LLN) with [I-D.ietf-6lo-backbone-router].

Beyond 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 [RFC8279] proposes an optimized technique to enable multicast in an LLN with a very limited requirement for routing state in the nodes.

Related requirements are:

Req2.1: The ND registration method SHOULD be extended so that the 6LR is instructed whether to advertise the Address of a 6LN 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 section 6.4 of [RFC6550], 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 Routing Registrar 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 EUI-64 address. At this point, the 6lo Working Group is extending the 6LoWPAN Header Compression (HC) [RFC6282] technique to other link types including 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 Bluetooth(R) Low Energy [RFC7668], and Power Line Communication (PLC) [I-D.hou-6lo-plc] Networks.

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 Neighbor 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 awake to answer a lookup from a node that uses IPv6 ND and may need a proxy. Additionally, the duty-cycled device may rely on the 6LBR to perform registration to the Routing Registrar.

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 SHOULD enable a Routing Registrar to operate as a proxy to defend the Registered Addresses on its behalf.

Req4.3: The registration mechanism SHOULD enable long sleep durations, on 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, spoofing the roles of the 6LR, 6LBR, and Routing Registrar 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 an 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 Routing Registrar 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 prevented.

Req5.3: 6LoWPAN ND security mechanisms SHOULD NOT lead to large 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 used.

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. Algorithm agility and support for large keys (e.g., 256-bit key sizes) is also desirable, following at Layer-3 the introduction of those capabilities at Layer-2.

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 6LN 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 to more hops.

B.7. Requirements Related to Operations and Management

Section 3.8 of "Architectural Principles of the Internet" [RFC1958] recommends to: "avoid options and parameters whenever possible. Any options and parameters should be configured or negotiated dynamically rather than manually". This is especially true in LLNs where the number of devices may be large and manual configuration is infeasible. Capabilities for a dynamic configuration of LLN devices can also be constrained by the network and power limitation.

A Network Administrator should be able to validate that the network is operating within capacity, and that in particular a 6LBR does not get overloaded with an excessive amount of registration, so the administrator can take actions such as adding a Backbone Link with additional 6LBRs and Routing Registrars to the network.

Related requirements are:

Req7.1: A management model SHOULD be provided that enables access to the 6LBR, monitor its usage vs. capacity, and alert in case of congestion. It is recommended that the 6LBR be reachable over a non-LLN link.

Req7.2: A management model SHOULD be provided that enables access to the 6LR and its capacity to host additional NCE. This management model SHOULD avoid polling individual 6LRs in a way that could disrupt the operation of the LLN.

Req7.3: Information on successful and failed registration SHOULD be provided, including information such as the ROVR of the 6LN, the Registered Address, the address of the 6LR, and the duration of the registration flow.

Req7.4: In case of a failed registration, information on the failure including the identification of the node that rejected the registration and the status in the EARO SHOULD be provided.

B.8. Matching Requirements with Specifications

I-drafts/RFCs addressing requirements

Requirement	Document
Req1.1	[I-D.ietf-6lo-backbone-router]
Req1.2	[RFC6775]
Req1.3	[RFC6775]
Req1.4	This RFC
Req2.1	This RFC
Req2.2	This RFC
Req2.3	
Req3.1	Technology Dependent
Req3.2	Technology Dependent
Req3.3	Technology Dependent
Req3.4	Technology Dependent
Req4.1	This RFC
Req4.2	This RFC
Req4.3	[RFC6775]
Req5.1	
Req5.2	[I-D.ietf-6lo-ap-nd]

Req5.3	
Req5.4	
Req5.5	[I-D.ietf-6lo-ap-nd]
Req5.6	[I-D.struik-lwip-curve-representations]
Req5.7	[I-D.ietf-6lo-ap-nd]
Req5.8	
Req5.9	[I-D.ietf-6lo-ap-nd]
Req6.1	This RFC
Req6.2	This RFC
Req7.1	
Req7.2	
Req7.3	
Req7.4	

Table 8: Work Addressing requirements

Authors' Addresses

Pascal Thubert (editor)
Cisco Systems, Inc
Building D (Regus) 45 Allee des Ormes
Mougins - Sophia Antipolis
France

Phone: +33 4 97 23 26 34
Email: pthubert@cisco.com

Erik Nordmark
Zededa
Santa Clara, CA
United States of America

Email: nordmark@sonic.net

Samita Chakrabarti
Verizon
San Jose, CA
United States of America

Email: samitac.ietf@gmail.com

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

Email: charliep@computer.org