

6lo
Internet-Draft
Intended status: Standards Track
Expires: January 9, 2020

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

Packet Delivery Deadline time in 6LoWPAN Routing Header
draft-ietf-6lo-deadline-time-05

Abstract

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

Status of This Memo

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

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

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

This Internet-Draft will expire on January 9, 2020.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1. Introduction	2
2. Terminology	3
3. 6LoRHE Generic Format	3
4. Deadline-6LoRHE	4
5. Deadline-6LoRHE Format	6
6. Deadline-6LoRHE in Three Network Scenarios	8
6.1. Scenario 1: Endpoints in the same DODAG (N1)	9
6.2. Scenario 2: Endpoints in Networks with Dissimilar L2 Technologies.	10
6.3. Scenario 3: Packet transmission across different DODAGs (N1 to N2).	11
7. IANA Considerations	12
8. Synchronization Aspects	13
9. Security Considerations	14
10. Acknowledgements	15
11. References	15
11.1. Normative References	15
11.2. Informative References	17
Appendix A. Changes from revision 04 to revision 05	18
Appendix B. Changes from revision 03 to revision 04	18
Appendix C. Changes from revision 02 to revision 03	19
Appendix D. Changes from revision 01 to revision 02	19
Appendix E. Changes between earlier versions	20
Authors' Addresses	20

1. Introduction

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

This document specifies a new type for the Elective 6LoWPAN Routing Header (6LoRHE), so that the deadline time (i.e., the time of latest acceptable delivery) of data packets can be included within the 6LoWPAN routing header. [RFC8138] specifies the 6LoWPAN Routing Header (6LoRH), compression schemes for RPL routing (source routing) operation [RFC6554], header compression of RPL Packet Information [RFC6553], and IP-in-IP encapsulation. This document also specifies handling of the deadline time when packets traverse between time-synchronized networks operating in different timezones or distinct reference clocks. Time synchronization techniques are outside the scope of this document. There are a number of standards available for this purpose, including IEEE 1588 [ieee-1588], IEEE 802.1AS [dot1AS-2011], IEEE 802.15.4-2015 TSCH [dot15-tsch], and more.

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

2. Terminology

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

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

3. 6LoRHE Generic Format

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




Figure 1: 6LoRHE format

- o Length: Length of the 6LoRHE expressed in bytes, excluding the first 2 bytes. This enables a node to skip a 6LoRHE if the Type is not recognized/supported.
- o Type (variable length): Type of the 6LoRHE (see Section 7)

4. Deadline-6LoRHE

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

The deadline field contains the value of the deadline time for the packet -- in other words, the time by which the application expects the packet to be delivered to the Receiver.

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

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

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

$$\text{new_network_origin_time} = \text{time_now_in_new_network} - \text{current_dly}$$

The following example illustrates these calculations when a packet travels between three networks, each in a different time zone. 'x' can be 1, 2 or 3. Suppose that the deadline time as measured in timezone 1 is 1050 and the origination time is 50. Suppose that the difference between TZ2 and TZ1 is 900, and the difference between TZ3 and TZ2 is 3600. In the figure, OT is the origination time as measured in the current timezone, and is equal to DT - OTD, that is, DT - 1000. Figure 2 uses the following abbreviations:

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

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

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

TZx : The time zone of network 'x'

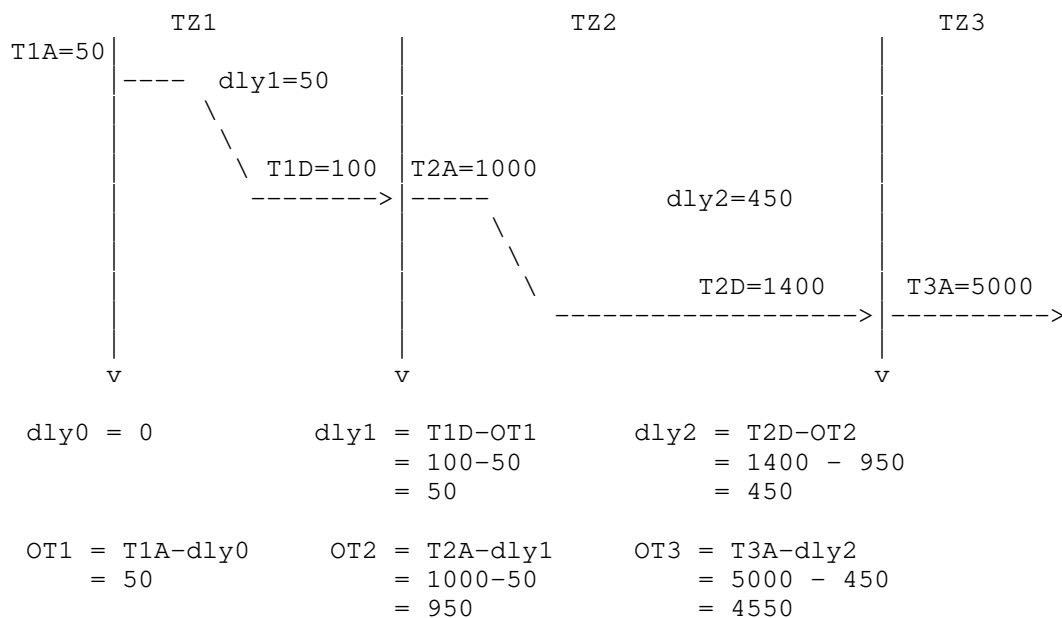


Figure 2: Destination Time Update example

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

5. Deadline-6LoRHE Format




Figure 3: Deadline-6LoRHE format

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

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

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

For the chosen time unit, a compressed time representation is available as follows. First, the application on the originating node has to determine how many time bits are needed to represent the difference between the time at which the packet is launched and the deadline time, including the representation of fractional time units. That number of bits (say, N_bits) determines DTL (the length of the Deadline Time (DT)) as follows:

$$DTL = (N_bits \bmod 4)$$

The number of bits determined by DTL allows counting any number of fractional time units in the range of interest determined by DT and the origination time OT. Denote this number of fractional time units to be Epoch_Range(DTL) (i.e., Epoch_Range is a function of DTL).

$$\text{Epoch_Range}(\text{DTL}) = (2^{(4 * (\text{DTL} + 1))})$$

Each point of time between OT and DT is represented by a time unit and a fractional time unit; in this section, this combined representation is called a rational time unit (RTU). 1 RTU measures the smallest fractional time that can be represented between two points of time in the epoch (i.e., within the range of interest).

DT - OT cannot exceed $2^{(4 * (\text{DTL} + 1))} = 16^{(\text{DTL} + 1)}$. A low value of DTL leads to a small Epoch_Range; if DTL = 0, there will only be 16 RTUs within the Epoch_Range (DTL) = 16^1 (for any time unit TU). The values that can be represented in the current epoch are in the range $[0, (\text{Epoch_Range}(\text{DTL}) - 1)]$. To minimize the required DTL,

wraparound is allowed but works naturally with the arithmetic modulo Epoch_Range.

By default, DTL determines t_0 in the chosen RTUs as follows:

$$t_0 = [\text{current_time} - (\text{current_time} \bmod \text{Epoch_Range}(\text{DTL}))].$$

Naturally, t_0 occurs at time 0 (or time 0.0000...) in the current epoch. The last possible origination time representable in the current epoch (counted in RTUs) is $t_{\text{last}} = (t_0 + (2^{(4*(\text{DTL}+1))-1}))$. In the RTUs chosen, the current epoch resides at the underlying time interval $[t_0, t_{\text{last}}]$. If $\text{DT} - \text{OT}$ is greater than $t_{\text{last}} - \text{OT}$, then wraparound within the Epoch_Range occurs naturally. In all cases, OT is represented by the value $(\text{OT} \bmod \text{Epoch_Range})$ and DT is represented by the value $(\text{DT} \bmod \text{Epoch_Range})$. All arithmetic is to be performed modulo $(\text{Epoch_Range}(\text{DTL}))$, yielding only positive values for $\text{DT} - \text{OT}$.

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

$$\begin{aligned} \text{deadline_time} &= \text{packet_origination_time} + \text{max_delay} \\ &= 0\text{x}D480 + 0\text{x}64 \text{ (Network ASNs)} \\ &= 0\text{x}D4E4 \text{ (Network ASNs)} \end{aligned}$$

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

$$\begin{aligned} \text{DTL} &= 3, \text{OTL} = 2, \text{TU} = 0\text{b}10, \text{BinaryPt} = 8, \text{DT} = 0\text{x}D4E4, \text{OTD} \\ &= 0\text{x}64 \end{aligned}$$

6. Deadline-6LoRHE in Three Network Scenarios

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




Figure 4: Intra-network Timezone Scenario

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

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

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




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

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

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

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




Figure 6: Packet transmission in Dissimilar L2 Technologies or Internet

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

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

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




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

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

```
current_time = ASN at LBR * slot_length_value

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

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

7. IANA Considerations

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

Elective 6LoRH Type	Value
Deadline-6LoRHE	TBD

Figure 8: Deadline-6LoRHE type

8. Synchronization Aspects

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

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

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

For example, suppose that $DTL = 0b0000$ and the DT bits are split evenly; then we can count up to 3.75 seconds by quarter-seconds.

If $DTL = 3$ and the DT bits are again split evenly, then we can count up to 256 seconds (in steps of $1/256$ of a second).

In all cases, t_0 is defined as specified in Section 5

$$t_0 = [\text{current_time} - (\text{current_time} \bmod (2^{4*(DTL+1)}))]$$

regardless of the choice of TU.

For TU = 0b00, the time units are seconds. With DTL == 15, and Binary Pt == 0, the epoch is (by default) January 1, 1900 at 00:00 UTC. The resolution is then $(2^{(-32)})$ seconds, which is the maximum possible. This time format wraps around every 2^{32} seconds, which is roughly 136 years.

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

9. Security Considerations

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

The protocol elements specified in this document are designed to work in controlled operational environments (e.g., industrial process control and automation). In order to avoid misuse of the deadline information that could potentially result in a Denial of Service (DoS) attack, proper functioning of this deadline time mechanism requires the provisioning and management of network resources for supporting traffic flows with deadlines, performance monitoring, and admission control policy enforcement. The network provisioning can be done either centrally or in a distributed fashion. For example, tracks in a 6tisch network could be established by a centralized PCE, as described in the 6tisch architecture [I-D.ietf-6tisch-architecture].

The Security Considerations of Detnet architecture [I-D.ietf-detnet-architecture] mostly apply to this document as well, as follows. To secure the request and control of resources allocated for tracks, authentication and authorization can be used for each device, and network controller devices. In the case of distributed control protocols, security is expected to be provided by the security properties of the protocols in use.

When deadline bearing flows are identified on a per-flow basis, which may provide attackers with additional information about the data flows, when compared to networks that do not include per-flow identification. The security implications of disclosing that additional information deserve consideration when implementing this deadline specification.

Because of the requirement of precise time synchronization, the accuracy, availability, and integrity of time synchronization is of critical importance. Extensive discussion of this topic can be found in [RFC7384].

10. Acknowledgements

The authors thank Pascal Thubert for suggesting the idea and encouraging the work. Thanks to Shwetha Bhandari's suggestions which were instrumental in extending the timing information to heterogeneous networks. The authors acknowledge the 6TiSCH WG members for their inputs on the mailing list. Special thanks to Jerry Daniel, Dan Frost (Routing Directorate) Charlie Kaufman (Security Directorate) Seema Kumar, Tal Mizrahi Avinash Mohan, Shalu Rajendran, Anita Varghese, and Dale Worley (Gen-ART review) for their support and valuable feedback.

11. References

11.1. Normative References

- [I-D.ietf-6tisch-terminology]
Palattella, M., Thubert, P., Watteyne, T., and Q. Wang,
"Terms Used in IPv6 over the TSCH mode of IEEE 802.15.4e",
draft-ietf-6tisch-terminology-10 (work in progress), March
2018.
- [I-D.ietf-detnet-architecture]
Finn, N., Thubert, P., Varga, B., and J. Farkas,
"Deterministic Networking Architecture", draft-ietf-
detnet-architecture-13 (work in progress), May 2019.
- [I-D.ietf-roll-routing-dispatch]
Thubert, P., Bormann, C., Toutain, L., and R. Cragie,
"6LoWPAN Routing Header", draft-ietf-roll-routing-
dispatch-05 (work in progress), October 2016.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", BCP 14, RFC 2119,
DOI 10.17487/RFC2119, March 1997,
<<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler,
"Transmission of IPv6 Packets over IEEE 802.15.4
Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007,
<<https://www.rfc-editor.org/info/rfc4944>>.

- [RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <<https://www.rfc-editor.org/info/rfc5905>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<https://www.rfc-editor.org/info/rfc6282>>.
- [RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <<https://www.rfc-editor.org/info/rfc6550>>.
- [RFC6553] Hui, J. and JP. Vasseur, "The Routing Protocol for Low-Power and Lossy Networks (RPL) Option for Carrying RPL Information in Data-Plane Datagrams", RFC 6553, DOI 10.17487/RFC6553, March 2012, <<https://www.rfc-editor.org/info/rfc6553>>.
- [RFC6554] Hui, J., Vasseur, JP., Culler, D., and V. Manral, "An IPv6 Routing Header for Source Routes with the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6554, DOI 10.17487/RFC6554, March 2012, <<https://www.rfc-editor.org/info/rfc6554>>.
- [RFC7384] Mizrahi, T., "Security Requirements of Time Protocols in Packet Switched Networks", RFC 7384, DOI 10.17487/RFC7384, October 2014, <<https://www.rfc-editor.org/info/rfc7384>>.
- [RFC7554] Watteyne, T., Ed., Palattella, M., and L. Grieco, "Using IEEE 802.15.4e Time-Slotted Channel Hopping (TSCH) in the Internet of Things (IoT): Problem Statement", RFC 7554, DOI 10.17487/RFC7554, May 2015, <<https://www.rfc-editor.org/info/rfc7554>>.
- [RFC8138] Thubert, P., Ed., Bormann, C., Toutain, L., and R. Cragie, "IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing Header", RFC 8138, DOI 10.17487/RFC8138, April 2017, <<https://www.rfc-editor.org/info/rfc8138>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

11.2. Informative References

[dot15-tsch]

"IEEE 802 Wireless", "IEEE Standard for Low-Rate Wireless Networks, Part 15.4, IEEE Std 802.15.4-2015", April 2016.

[dot1AS-2011]

"IEEE Standards", "IEEE Standard for Local and Metropolitan Area Networks - Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks", March 2011.

[dotBLEMesh]

Leonardi, L., Pattim, G., and L. Lo Bello, "Multi-Hop Real-Time Communications Over Bluetooth Low Energy Industrial Wireless Mesh Networks", IEEE Access Vol 6, 26505-26519, May 2018.

[dotWi-SUN]

Harada, H., Mizutani, K., Fujiwara, J., Mochizuki, K., Obata, K., and R. Okumura, "IEEE 802.15.4g Based Wi-SUN Communication Systems", IEICE Transactions on Communications volume E100.B, Jan 2017.

[I-D.ietf-6lo-backbone-router]

Thubert, P., Perkins, C., and E. Levy-Abegnoli, "IPv6 Backbone Router", draft-ietf-6lo-backbone-router-11 (work in progress), February 2019.

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

Gomez, C., Darroudi, S., Savolainen, T., and M. Spoerk, "IPv6 Mesh over BLUETOOTH(R) Low Energy using IPSP", draft-ietf-6lo-blemesh-05 (work in progress), March 2019.

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

Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", draft-ietf-6tisch-architecture-24 (work in progress), July 2019.

[I-D.ietf-detnet-use-cases]

Grossman, E., "Deterministic Networking Use Cases", draft-ietf-detnet-use-cases-20 (work in progress), December 2018.

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

Brockners, F., Bhandari, S., Pignataro, C., Gredler, H., Leddy, J., Youell, S., Mizrahi, T., Mozes, D., Lapukhov, P., Chang, R., daniel.bernier@bell.ca, d., and J. Lemon, "Data Fields for In-situ OAM", draft-ietf-ippm-ioam-data-06 (work in progress), July 2019.

[I-D.ietf-roll-useofrplinfo]

Robles, I., Richardson, M., and P. Thubert, "Using RPL Option Type, Routing Header for Source Routes and IPv6-in-IPv6 encapsulation in the RPL Data Plane", draft-ietf-roll-useofrplinfo-31 (work in progress), July 2019.

[ieee-1588]

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

[Wi-SUN_PHY]

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

Appendix A. Changes from revision 04 to revision 05

This section lists the changes between draft-ietf-6lo-deadline-time revisions ...-04.txt and ...-05.txt.

- o Included additional relevant material in Security Considerations regarding expected deployment scenarios and the effect of disclosing additional information during the travel of a packet.
- o Reworked the specification for using time ranges shorter than the maximum allowed by the choice of TU, so that fewer bits are needed to represent DT and OT.
- o Revised the figures and examples to use new parameters
- o Reordered the field definitions for the Deadline-6LoRHE.
- o Responded to numerous reviewer comments to improve terminology and editorial consistency.

Appendix B. Changes from revision 03 to revision 04

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

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

Appendix C. Changes from revision 02 to revision 03

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

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

Appendix D. Changes from revision 01 to revision 02

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

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

- o Updated bibliographic citations.
- o Alphabetized contributor names.

Appendix E. Changes between earlier versions

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

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

Authors' Addresses

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

Email: lijo@cdac.in

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

Email: satishnaidu80@gmail.com

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

Email: anand@ece.iisc.ernet.in

Malati Hegde
Indian Institute of Science
Bangalore 560012
India

Email: malati@ece.iisc.ernet.in

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

Email: charliep@computer.org

6lo
Internet-Draft
Updates: 4944 (if approved)
Intended status: Standards Track
Expires: December 13, 2019

P. Thubert, Ed.
Cisco Systems
June 11, 2019

6LoWPAN Selective Fragment Recovery
draft-ietf-6lo-fragment-recovery-04

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 13, 2019.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1. Introduction	2
2. Terminology	4
2.1. BCP 14	4
2.2. References	4
2.3. 6LoWPAN Acronyms	4
2.4. Referenced Work	4
2.5. New Terms	5
3. Updating RFC 4944	6
4. Updating draft-ietf-6lo-minimal-fragment	6
4.1. Slack in the First Fragment	7
4.2. Gap between frames	7
4.3. Modifying the First Fragment	7
5. New Dispatch types and headers	8
5.1. Recoverable Fragment Dispatch type and Header	9
5.2. RFRAG Acknowledgment Dispatch type and Header	11
6. Fragments Recovery	13
6.1. Forwarding Fragments	15
6.1.1. Upon the first fragment	15
6.1.2. Upon the next fragments	15
6.2. Upon the RFRAG Acknowledgments	16
6.3. Aborting the Transmission of a Fragmented Packet	17
7. Management Considerations	17
7.1. Protocol Parameters	17
7.2. Observing the network	18
8. Security Considerations	19
9. IANA Considerations	19
10. Acknowledgments	19
11. References	19
11.1. Normative References	19
11.2. Informative References	20
Appendix A. Rationale	22
Appendix B. Requirements	24
Appendix C. Considerations On Flow Control	24
Author's Address	26

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 a payload of 74 bytes or more, 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 the firmware upgrade of the LLN nodes or the 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 10 kilobytes 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.ietf-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 issues of resources locked and wasted transmissions due to the loss of a fragment. [RFC4944] does not define a mechanism to first discover a fragment loss, 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. Constrained memory resources are blocked on the receiver until the receiver times out.

That problem is exacerbated when forwarding fragments over multiple hops since a loss at an intermediate hop will not be discovered by either the source or the destination, and the source will keep on sending fragments, wasting even more resources in the network and possibly contributing to the condition that caused the loss to no avail since the datagram cannot arrive in its entirety. RFC 4944 is also 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 potentially the retransmission of the full datagram.

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

2. Terminology

2.1. BCP 14

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

2.2. References

In this document, readers will encounter terms and concepts that are discussed in "Problem Statement and Requirements for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing" [RFC6606]

2.3. 6LoWPAN Acronyms

This document uses the following acronyms:

6BBR: 6LoWPAN Backbone Router

6LBR: 6LoWPAN Border Router

6LN: 6LoWPAN Node

6LR: 6LoWPAN Router

LLN: Low-Power and Lossy Network

2.4. Referenced Work

Past experience with fragmentation has shown that misassociated 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.ietf-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 MPLS. This specification reuses that technique with slightly modified controls.

2.5. New Terms

This specification uses the following terms:

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

Compressed Form This specification uses the generic term Compressed Form to refer to the format of a datagram after the action of [RFC6282] and possibly [RFC8138] for RPL [RFC6550] artifacts.

`datagram_size`: The size of the datagram in its Compressed Form before it is fragmented. The `datagram_size` is expressed in a unit that depends on the MAC layer technology, by default a byte.

`fragment_offset`: The offset of a particular fragment of a datagram in its Compressed Form. The `fragment_offset` is expressed in a unit that depends on the MAC layer technology and is by default a byte.

`datagram_tag`: An identifier of a datagram that is locally unique to the Layer-2 sender. Associated with the MAC address of the

sender, this becomes a globally unique identifier for the datagram.

RFRAG: Recoverable Fragment

RFRAG-ACK: Recoverable Fragment Acknowledgement

RFRAG Acknowledgment Request: An RFRAG with the Acknowledgement Request flag ('X' flag) set.

All 0's: Refers to a bitmap with all bits set to zero.

All 1's: Refers to a bitmap with all bits set to one.

3. Updating RFC 4944

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

[RFC8138] allows to modify the size of a packet en-route by removing the consumed hops in a compressed Routing Header. It results that `fragment_offset` and `datagram_size` (see Section 2.5) must also be modified en-route, which is difficult to do in the uncompressed form. This specification expresses those fields in the Compressed Form and allows to modify them en-route (see Section 4.3) easily.

Note that consistently with Section 2 of [RFC6282] for the fragmentation mechanism described in Section 5.3 of [RFC4944], any header that cannot fit within the first fragment MUST NOT be compressed when using the fragmentation mechanism described in this specification.

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

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

4.1. Slack in the First Fragment

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

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

4.2. Gap between frames

This specification introduces a concept of Inter-Frame Gap, which is a configurable interval of time between transmissions to a same next hop. In the case of half duplex interfaces, this `InterFrameGap` ensures that the next hop has progressed the previous frame and is capable of receiving the next one.

In the case of a mesh operating at a single frequency with omnidirectional antennas, a larger `InterFrameGap` is required to protect the frame against hidden terminal collisions with the previous frame of a same flow that is still progressing along a common path.

The Inter-Frame Gap is useful even for unfragmented datagrams, but it becomes a necessity for fragments that are typically generated in a fast sequence and are all sent over the exact same path.

4.3. Modifying the First Fragment

The compression of the Hop Limit, of the source and destination addresses in the IPv6 Header, and of the Routing Header, may change en-route in a Route-Over mesh LLN. If the size of the first fragment is modified, then the intermediate node MUST adapt the `datagram_size` to reflect that difference.

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

5. New Dispatch types and headers

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

This specification provides a technique that is derived from MPLS to forward individual fragments across a 6LoWPAN route-over mesh without reassembly at each hop. The `datagram_tag` is used as a label; it is locally unique to the node that owns 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 chosen `datagram_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 2 new Dispatch types, for Recoverable Fragment (RFRAG) and for the RFRAG Acknowledgment back.

(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 10100x	RFRAG - Recoverable Fragment
11 10101x	RFRAG-ACK - RFRAG Acknowledgment

Figure 1: Additional Dispatch Value Bit Patterns

In the following sections, a "`datagram_tag`" extends the semantics defined in [RFC4944] Section 5.3. "Fragmentation Type and Header". The `datagram_tag` is a locally unique identifier for the datagram from the perspective of the sender. This means that the `datagram_tag` identifies a datagram uniquely in the network when associated with

the source of the datagram. As the datagram gets forwarded, the source changes and the `datagram_tag` must be swapped as detailed in [I-D.ietf-6lo-minimal-fragment].

5.1. Recoverable Fragment Dispatch type and Header

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

The format of the fragment header is shown in Figure 2. It is the same for all fragments. The format has a length and an offset, as well as a sequence field. This would be redundant if the offset was computed as the product of the sequence by the length, but this is not the case. The position of a fragment in the reassembly buffer is neither correlated with the value of the sequence field nor with the order in which the fragments are received. This enables out-of-sequence and overlapping fragments, e.g., a fragment 5 that is retried as smaller fragments 5, 13 and 14 due to a change of MTU.

There is no requirement on the receiver to check for contiguity of the received fragments, and the sender MUST ensure that when all fragments are acknowledged, then the datagram is fully received. This may be useful in particular in the case where the MTU changes and a fragment sequence is retried with a smaller `fragment_size`, the remainder of the original fragment being retried with new sequence values.

The first fragment is recognized by a sequence of 0; it carries its `fragment_size` and the `datagram_size` of the compressed packet before it is fragmented, 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-Request: 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: 10 bit unsigned integer; the size of this fragment in a unit that depends on the MAC layer technology. By default, that unit is the octet which allows fragments up to 512 bytes. For IEEE Std. 802.15.4, the unit is octet, and the maximum fragment size, when it 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.
- datagram_tag: 16 bits; an identifier of the datagram that is locally unique to the sender.
- Sequence: 5 bit unsigned integer; the sequence number of the fragment in the acknowledgement bitmap. Fragments are numbered [0..N] where N is in [0..31]. A Sequence of 0 indicates the first fragment in a datagram, but non-zero values are not indicative of the position in the reassembly buffer.
- Fragment_offset: 16 bit unsigned integer;
- * When the Fragment_offset is set to a non-0 value, its semantics depend on the value of the Sequence field.
 - + For a first fragment (i.e. with a Sequence of 0), this field indicates the datagram_size of the compressed datagram, to help the receiver allocate an adapted buffer for the reception and reassembly operations. The fragment may be stored for local reassembly. Alternatively, it may be routed based on the destination IPv6 address. In that case, a VRB state must be installed as described in Section 6.1.1.

- + When the Sequence is not 0, this field indicates the offset of the fragment in the Compressed Form of the datagram. The fragment may be added to a local reassembly buffer or forwarded based on an existing VRB as described in Section 6.1.2.
- * A Fragment_offset that is set to a value of 0 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 6.1.2, 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 6.1.1 but no state is conserved afterwards. In that case, the session if it exists is aborted and the packet is also forwarded in an attempt to clean up the next hops as along the path indicated by the IPv6 header (possibly including a routing header).

If the fragment cannot be forwarded or routed, then an abort RFRAG-ACK is sent back to the source as described in Section 6.1.2.

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:

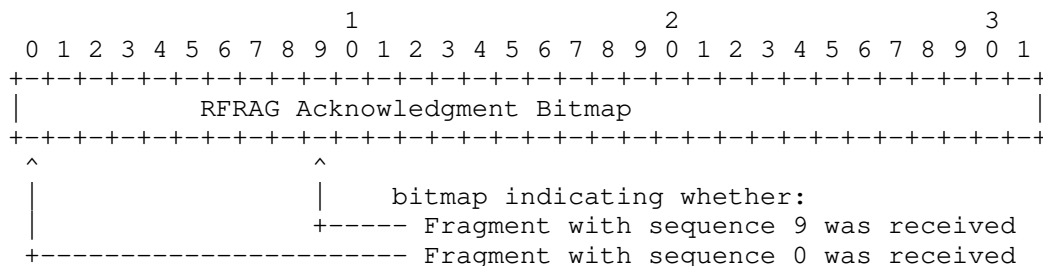


Figure 3: RFRAG Acknowledgment bitmap encoding

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 lost and must be retried.




Figure 4: Example RFRAG Acknowledgment Bitmap

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




Figure 5: RFRAG Acknowledgment Dispatch type and Header

E: 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. More in Appendix C.

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 header RFRAG is used to transport a fragment and optionally request an RFRAG Acknowledgment that will confirm the good reception of one or more fragments. An RFRAG Acknowledgment is carried as a standalone fragment header (i.e. with no 6LoWPAN payload) in a message that is propagated back to the 6LoWPAN endpoint that was the originator of the fragments. To achieve this, each hop that performed an MPLS-like operation on fragments reverses that operation for the RFRAG_ACK by sending a frame from the next hop to the previous hop as known by its MAC address in the VRB. The *datagram_tag* in the RFRAG_ACK is unique to the receiver and is enough information for an intermediate hop to locate the VRB that contains the *datagram_tag* used by the previous hop and the Layer-2 information associated to it (interface and MAC address).

The 6LoWPAN endpoint that fragments the packets at 6LoWPAN level (the sender) also controls the amount of acknowledgments by setting the Ack-Request flag in the RFRAG packets. The sender may set the Ack-Request 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. The maximum number of outstanding fragments is the Window-Size. It is configurable and may vary in case of ECN notification. When the 6LoWPAN endpoint that reassembles the packets at 6LoWPAN level (the receiver) receives a fragment with the Ack-Request flag set, it MUST send an RFRAG Acknowledgment back to the originator to confirm reception of all the fragments it has received so far.

The Ack-Request ('X') set in an RFRAG marks the end of a window. This flag SHOULD be set on the last fragment to protect the datagram, and it MAY be set in any intermediate fragment for the purpose of flow control. This ARQ process MUST be protected by a timer, and the fragment that carries the 'X' flag MAY be retried upon time out a configurable amount of times (see Section 7.1). Upon exhaustion of the retries the sender may either abort the transmission of the datagram or retry the datagram from the first fragment with an 'X' flag set in order to reestablish a path and discover which fragments were received over the old path in the acknowledgment bitmap. 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-Request bit.

The RFRAG Acknowledgment can optionally carry an ECN indication for flow control (see Appendix C). The receiver of a fragment with the 'E' (ECN) flag set MUST echo that information by setting the 'E' (ECN) flag in the next RFRAG Acknowledgment.

The sender transfers a controlled number of fragments and MAY flag the last fragment of a window with an RFRAG Acknowledgment Request. The receiver MUST acknowledge a fragment with the acknowledgment request bit set. If any fragment immediately preceding an acknowledgment request is still missing, the receiver MAY intentionally delay its acknowledgment to allow in-transit fragments to arrive. Because it might defeat the round trip delay computation, delaying the acknowledgment 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 aborted 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 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]

6.1. 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. It inherits operations like `datagram_tag` Switching and using a timer to clean the VRB when the traffic dries up. In more details, 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, the next fragments can only follow the path that was set up by the first fragment and cannot follow an alternate route. 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.

6.1.1. 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.ietf-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.

6.1.2. 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.ietf-6lo-minimal-fragment].

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

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

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

6.2. 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.ietf-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 that the fragment was entirely received (FULL bitmap), it arms a short timer, and upon timeout, the VRB and all the associated state are destroyed. Until the timer elapses, 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 forwards the fragment according to the state in the VRB.

This specification does not provide a method to discover the number of hops or the minimal value of MTU along those hops. But should the minimal MTU decrease, it is possible to retry a long fragment (say sequence of 5) with first a shorter fragment of the same sequence (5 again) and then one or more other fragments with a sequence that was not used before (e.g., 13 and 14). Note that Path MTU Discovery is out of scope for this document.

6.3. Aborting the Transmission of a Fragmented Packet

A reset is signaled on the forward path with a pseudo fragment that has the `fragment_offset`, sequence and `fragment_size` all set to 0, and no data.

When the sender or a router on the way decides that a packet should be dropped and the fragmentation process aborted, it generates a reset pseudo fragment and forwards it down the fragment path.

Each router next along the path the way forwards the pseudo fragment based on the VRB state. If an acknowledgment is not requested, the VRB and all associated state are destroyed.

Upon reception of the pseudo fragment, the receiver cleans up all resources for the packet associated to the `datagram_tag`. If an acknowledgment is requested, the receiver responds with a NULL bitmap.

The other way around, the receiver might need to abort 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. Management Considerations

7.1. Protocol Parameters

There is no particular configuration on the receiver, as echoing ECN is always on. The configuration only applies to the sender, which is in control of the transmission. The management system SHOULD be capable of providing the parameters below:

MinFragmentSize: The `MinFragmentSize` is the minimum value for the `Fragment_Size`.

OptFragmentSize: The `MinFragmentSize` is the value for the `Fragment_Size` that the sender should use to start with.

MaxFragmentSize: The `MaxFragmentSize` is the maximum value for the `Fragment_Size`. It MUST be lower than the minimum MTU along the path. A large value augments the chances of buffer bloat and transmission loss. The value MUST be less than 512 if the unit that is defined for the PHY layer is the octet.

UseECN: Indicates whether the sender should react to ECN. When the sender reacts to ECN the Window_Size will vary between MinWindowSize and MaxWindowSize.

MinWindowSize: The minimum value of Window_Size that the sender can use.

OptWindowSize: The OptWindowSize is the value for the Window_Size that the sender should use to start with.

MaxWindowSize: The maximum value of Window_Size that the sender can use. The value MUST be less than 32.

InterFrameGap: Indicates a minimum amount of time between transmissions. All packets to a same destination, and in particular fragments, may be subject to receive while transmitting and hidden terminal collisions with the next or the previous transmission as the fragments progress along a same path. The InterFrameGap protects the propagation of one transmission before the next one is triggered and creates a duty cycle that controls the ratio of air time and memory in intermediate nodes that a particular datagram will use.

MinARQTimeout: The maximum amount of time a node should wait for an RFRAG Acknowledgment before it takes a next action.

OptARQTimeout: The starting point of the value of the amount that a sender should wait for an RFRAG Acknowledgment before it takes a next action.

MaxARQTimeout: The maximum amount of time a node should wait for an RFRAG Acknowledgment before it takes a next action.

MaxFragRetries: The maximum number of retries for a particular Fragment.

MaxDatagramRetries: The maximum number of retries from scratch for a particular Datagram.

7.2. Observing the network

The management system should monitor the amount of retries and of ECN settings that can be observed from the perspective of the both the sender and the receiver, and may tune the optimum size of Fragment_Size and of the Window_Size, OptWindowSize and OptWindowSize respectively, at the sender. The values should be bounded by the expected number of hops and reduced beyond that when the number of datagrams that can traverse an intermediate point may exceed its

capacity and cause a congestion loss. The InterFrameGap is another tool that can be used to increase the spacing between fragments of a same datagram and reduce the ratio of time when a particular intermediate node holds a fragment of that datagram.

8. Security Considerations

The considerations in the Security section of [I-D.ietf-core-cocoa] apply equally to this specification.

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 Michel Veillette, Dario Tedeschi, Laurent Toutain, Carles Gomez Montenegro, 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.ietf-6lo-minimal-fragment] Watteyne, T., Bormann, C., and P. Thubert, "LLN Minimal Fragment Forwarding", draft-ietf-6lo-minimal-fragment-01 (work in progress), March 2019.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<https://www.rfc-editor.org/info/rfc4944>>.

- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<https://www.rfc-editor.org/info/rfc6282>>.
- [RFC6554] Hui, J., Vasseur, JP., Culler, D., and V. Manral, "An IPv6 Routing Header for Source Routes with the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6554, DOI 10.17487/RFC6554, March 2012, <<https://www.rfc-editor.org/info/rfc6554>>.
- [RFC8025] Thubert, P., Ed. and R. Cragie, "IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Paging Dispatch", RFC 8025, DOI 10.17487/RFC8025, November 2016, <<https://www.rfc-editor.org/info/rfc8025>>.
- [RFC8138] Thubert, P., Ed., Bormann, C., Toutain, L., and R. Cragie, "IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing Header", RFC 8138, DOI 10.17487/RFC8138, April 2017, <<https://www.rfc-editor.org/info/rfc8138>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

11.2. Informative References

- [I-D.ietf-6tisch-architecture]
Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", draft-ietf-6tisch-architecture-20 (work in progress), March 2019.
- [I-D.ietf-core-cocoa]
Bormann, C., Betzler, A., Gomez, C., and I. Demirkol, "CoAP Simple Congestion Control/Advanced", draft-ietf-core-cocoa-03 (work in progress), February 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>>.
- [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>>.
- [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 would need to be 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 Allée des Ormes - BP1200
MOUGINS - Sophia Antipolis 06254
FRANCE

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

6lo
Internet-Draft
Intended status: Informational
Expires: December 26, 2019

T. Watteyne, Ed.
Analog Devices
C. Bormann
Universitaet Bremen TZI
P. Thubert
Cisco
June 24, 2019

LLN Minimal Fragment Forwarding
draft-ietf-6lo-minimal-fragment-02

Abstract

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

Status of This Memo

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

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

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

This Internet-Draft will expire on December 26, 2019.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents

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

Table of Contents

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

1. Overview of 6LoWPAN Fragmentation

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

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




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

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

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

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

Conceptually, a reassembly buffer for 6LoWPAN contains:

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

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

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

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

2. Limits of Per-Hop Fragmentation and Reassembly

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

2.1. Latency

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

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

2.2. Memory Management and Reliability

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

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




Figure 2: Illustrating the Memory Management Issue.

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

3. Virtual Reassembly Buffer (VRB) Implementation

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

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

There are, however, limits:

Non-zero Packet Drop Probability: The abstract data in a VRB table entry contains at a minimum the MAC address of the predecessor and that of the successor, the `datagram_tag` used by the predecessor and the local `datagram_tag` that this node will swap with it. The VRB may need to store a few octets from the last fragment that may not have fit within MTU and that will be prepended to the next fragment. This yields a small footprint that is 2 orders of magnitude smaller compared to needing a 1280-byte reassembly buffer for each packet. Yet, the size of the VRB table necessarily remains finite. In the extreme case where a node is required to concurrently forward more packets than it has entries in its VRB table, packets are dropped.

No Fragment Recovery: There is no mechanism in VRB for the node that reassembles a packet to request a single missing fragment. Dropping a fragment requires the whole packet to be resent. This causes unnecessary traffic, as fragments are forwarded even when the destination node can never construct the original IPv6 packet.

No Per-Fragment Routing: All subsequent fragments follow the same sequence of hops from the source to the destination node as the first fragment, because the IP header is required to route the fragment and is only present in the first fragment. A side effect is that the first fragment must always be forwarded first.

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

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

4. Security Considerations

An attacker can perform a Denial-of-Service (DoS) attack on a node implementing VRB by generating a large number of bogus "fragment 1" fragments without sending subsequent fragments. This causes the VRB table to fill up. Note that the VRB does not need to remember the full datagram as received so far but only possibly a few octets from the last fragment that could not fit in it. It is expected that an implementation protects itself to keep the number of VRBs within capacity, and that old VRBs are protected by a timer of a reasonable duration for the technology and destroyed upon timeout.

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

5. IANA Considerations

No requests to IANA are made by this document.

6. Acknowledgments

The authors would like to thank Yasuyuki Tanaka, for his in-depth review of this document. Also many thanks to Georgies Papadopoulos and Dominique Barthel for their own reviews.

7. Informative References

- [ARTICLE] Tanaka, Y., Minet, P., and T. Watteyne, "6LoWPAN Fragment Forwarding", IEEE Communications Standards Magazine , 2019.
- [BOOK] Shelby, Z. and C. Bormann, "6LoWPAN", John Wiley & Sons, Ltd monograph, DOI 10.1002/9780470686218, November 2009.
- [I-D.ietf-6lo-fragment-recovery]
Thubert, P., "6LoWPAN Selective Fragment Recovery", draft-ietf-6lo-fragment-recovery-04 (work in progress), June 2019.
- [I-D.ietf-lwig-6lowpan-virtual-reassembly]
Bormann, C. and T. Watteyne, "Virtual reassembly buffers in 6LoWPAN", draft-ietf-lwig-6lowpan-virtual-reassembly-01 (work in progress), March 2019.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<https://www.rfc-editor.org/info/rfc4944>>.

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

Authors' Addresses

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

Email: thomas.watteyne@analog.com

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

Email: cabo@tzi.org

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

Email: pthubert@cisco.com

6Lo Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 9, 2020

Y. Choi, Ed.
Y-G. Hong
ETRI
J-S. Youn
Dongueui Univ
D-K. Kim
KNU
J-H. Choi
Samsung Electronics Co.,
July 8, 2019

Transmission of IPv6 Packets over Near Field Communication
draft-ietf-6lo-nfc-15

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 January 9, 2020.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

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

1. Introduction

NFC is a set of short-range wireless technologies, typically requiring a distance of 10 cm or less. NFC operates at 13.56 MHz on ISO/IEC 18000-3 air interface and at rates ranging from 106 kbit/s to

424 kbit/s [ECMA-340]. NFC always involves an initiator and a target; the initiator actively generates an RF field that can power a passive target. This enables NFC targets to take very simple form factors such as tags, stickers, key fobs, or cards that do not require batteries. NFC peer-to-peer communication is possible, provided both devices are powered. NFC builds upon RFID systems by allowing two-way communication between endpoints. At the time of this writing, it has been used in devices such as mobile phones, running Android operating system, named with a feature called "Android Beam". It is expected for the other mobile phones, running other operating systems (e.g., iOS, etc.) to be equipped with NFC technology in the near future.

Considering exponential growth in the number of heterogeneous air interface technologies, NFC has been widely used like Bluetooth Low Energy (BT-LE), Wi-Fi, and so on. Each of the heterogeneous air interface technologies has its own characteristics, which cannot be covered by the other technologies, so various kinds of air interface technologies would co-exist together. NFC can provide secured communications with its short transmission range.

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

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

2. Conventions and Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "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

NFC enables simple and two-way interaction between two devices, allowing consumers to perform contactless transactions, access digital content, and connect electronic devices with a single touch. NFC complements many popular consumer level wireless technologies, by utilizing the key elements in existing standards for contactless card

technology (ISO/IEC 14443 A&B and JIS-X 6319-4). NFC can be compatible with existing contactless card infrastructure and it enables a consumer to utilize one device across different systems.

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

3.1. Peer-to-peer Mode of NFC

NFC-enabled devices are unique in that they can support three modes of operation: card emulation, peer-to-peer, and reader/writer. Only peer-to-peer in the three modes enables two NFC-enabled devices to communicate with each other to exchange information and share files, so that users of NFC-enabled devices can quickly share contact information and other files with a touch. Therefore, the peer mode is used for ipv6-over-nfc. In addition, NFC-enabled devices can securely send IPv6 packets in wireless range when an NFC-enabled gateway is linked to the Internet.

3.2. Protocol Stacks of NFC

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

For data communication in IPv6 over NFC, an IPv6 packet MUST be passed down to LLCP of NFC and transported to an Information (I) and an Unnumbered Information (UI) Field in Protocol Data Unit (PDU) of LLCP of the NFC-enabled peer device. LLCP does not support fragmentation and reassembly. For IPv6 addressing or address configuration, LLCP MUST provide related information, such as link layer addresses, to its upper layer. The LLCP to IPv6 protocol binding MUST transfer the SSAP and DSAP value to the IPv6 over NFC protocol. SSAP stands for Source Service Access Point, which is a 6-bit value meaning a kind of Logical Link Control (LLC) address, while DSAP means an LLC address of the destination NFC-enabled device.

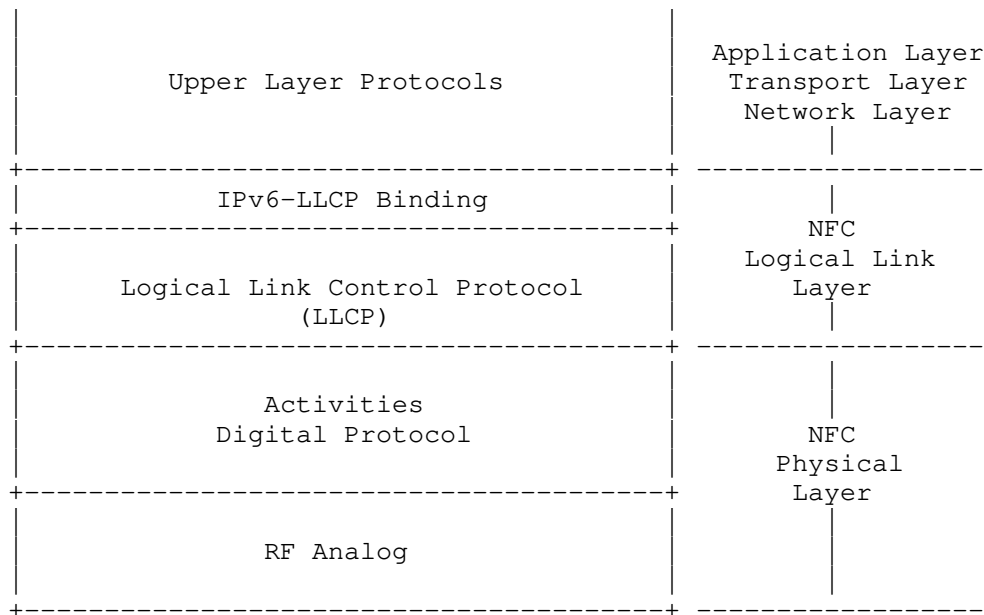


Figure 1: Protocol Stacks of NFC

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

3.3. NFC-enabled Device Addressing

According to NFC Logical Link Control Protocol 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. The several service access points can be installed on a NFC device. However, the SSAP and DSAP can be used as identifiers for NFC link connections with the IPv6 over NFC adaptation layer. Therefore, the SSAP can be used to generate an IPv6 interface identifier. Address values between 00h and 0Fh of SSAP and DSAP are reserved for identifying the well-known service access points, which are defined in the NFC Forum Assigned

Numbers Register. Address values between 10h and 1Fh 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, an IPv6 packet MUST be passed down to LLC of NFC and transported to an Unnumbered Information Protocol Data Unit (UI PDU) and an Information Field in Protocol Data Unit (I PDU) of LLC of the NFC-enabled peer device.

The information field of an I PDU 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 is announce a larger MIU for a data link connection by transmitting an MIUX extension parameter within the information field. If no MIUX parameter is transmitted, the MIU value is 128 bytes. Otherwise, the MTU size in NFC LLC 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. Each of TLV Type and TLV Length field is 1 byte, and TLV Value field is 2 bytes.

0	0	1	2	3				
0	8	6	2	1				
+-----+-----+-----+-----+								
	Type		Length				Value	
+-----+-----+-----+-----+								
	00000010		00000010		1011		0x0~0x7FF	
+-----+-----+-----+-----+								

Figure 2: Example of MIUX Parameter TLV

When the MIUX parameter is encoded as a TLV option, 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. A maximum value of the TLV Value field can be 0x7FF, and a maximum size of the MTU in

NFC LLCP is 2176 bytes including the 128 byte default of MIU. This value MUST be 0x480 to cover MTU of IPV6 if FAR is not used in IPV6 over NFC.

4. Specification of IPv6 over NFC

NFC technology also has considerations and requirements owing to low power consumption and allowed protocol overhead. 6LoWPAN standards [RFC4944], [RFC6775], and [RFC6282] provide useful functionality for reducing overhead which can be applied to NFC. This functionality consists of link-local IPv6 addresses and stateless IPv6 address auto-configuration (see Section 4.3), Neighbor Discovery (see Section 4.5) and header compression (see Section 4.7).

4.1. Protocol Stacks

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




Figure 3: Protocol Stacks for IPv6 over NFC

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

4.2. Link Model

In the case of BT-LE, the Logical Link Control and Adaptation Protocol (L2CAP) supports fragmentation and reassembly (FAR) functionality; therefore, the adaptation layer for IPv6 over BT-LE does not have to conduct the FAR procedure. The NFC LLCP, in contrast, does not support the FAR functionality, so IPv6 over NFC needs to consider the FAR functionality, defined in [RFC4944]. However, the MTU on an NFC link can be configured in a connection

procedure and extended enough to fit the MTU of IPv6 packet (see Section 4.8).

This document does NOT RECOMMEND using FAR over NFC link. In addition, the implementation for this specification MUST use MIUX extension to communicate the MTU of the link to the peer as defined in Section 3.4.

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

4.3. Stateless Address Autoconfiguration

An NFC-enabled device (i.e., 6LN) performs stateless address autoconfiguration as per [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 algorithm, F() with input parameters. One of the parameters is Net_IFace, and NFC Link Layer address (i.e., SSAP) is a source of the NetIFace parameter. The 6-bit address of SSAP of NFC is easy and short to be targeted by attacks of third party (e.g., address scanning). The F() 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.

4.4. 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

The tool for a 6LBR to obtain an IPv6 prefix for numbering the NFC network can be accomplished via DHCPv6 Prefix Delegation ([RFC3633]). The "Interface Identifier" is used the secured and stable IIDs for NFC-enabled devices.

4.5. Neighbor Discovery

Neighbor Discovery Optimization for 6LoWPANs ([RFC6775]) describes the neighbor discovery approach in several 6LoWPAN topologies, such as mesh topology. NFC does not support a complicated mesh topology but only a simple multi-hop network topology or directly connected peer-to-peer network. Therefore, the following aspects of RFC 6775 are applicable to NFC:

- o When an NFC-enabled device (6LN) is directly connected to a NFC-enabled 6LBR, an NFC 6LN MUST register its address with the 6LBR[RFC4944] by sending a Neighbor Solicitation (NS) message with the Address Registration Option (ARO) 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.
- o When two or more NFC 6LNs[RFC4944] (or 6LRs) 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 (e.g., isolated network). In a case of multi-hops, all of 6LNs, which have two or more connections with

different neighbors, is a router for 6LR/6LBR. In a case that they meet within a single hop and they have the same properties, any of them can be a router.

- o For sending Router Solicitations and processing Router Advertisements, the NFC 6LNs MUST follow Sections 5.3 and 5.4 of [RFC6775].
- o When a NFC device becomes a 6LR or a 6LBR, the NFC device MUST follow Section 6 and 7 of [RFC6775].

4.6. Dispatch Header

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




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

The dispatch value 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.

4.7. 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

4.8. Fragmentation and Reassembly Considerations

IPv6-over-NFC MUST NOT use fragmentation and reassembly (FAR) 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 fit the MTU of IPv6 Packet. The MIUX value is 0x480 in order to fit the MTU (1280 bytes) of a IPv6 packet if NFC devices support extension of the MTU. However, if the NFC device does not support extension, IPv6-over-NFC uses FAR with the default MTU (128 bytes), as defined in [RFC4944].

4.9. 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 be isolated and connected to the Internet.

5.1. NFC-enabled Device Connected to the Internet

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

Figure 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. If there is any laptop computers close to a user, it will become a 6LBR. Additionally, when the user mounts an NFC-enabled air interface adapter (e.g., portable NFC dongle) on the close laptop PC, the user's NFC-enabled device (6LN) can communicate with the laptop PC (6LBR) within 10 cm distance.




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

Two or more 6LNs are 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 transiently be a simple isolated network as shown in the Figure 11.




Figure 11: Isolated NFC-enabled device network

In mobile phone markets, applications are designed and made by user developers. They may image interesting applications, where three or

more mobile phones touch or attach each other to work achieve a common objective. In an isolated NFC-enabled device network, when two or more 6LRs are connected with each other, and then they are acting like routers, the 6LR MUST ensure address collisions do not occur.

6. IANA Considerations

There are no IANA considerations related to this document.

7. Security Considerations

This document does not RECOMMEND sending NFC packets over the Internet or any unsecured network.

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 the every connection is made with different address of NFC link with an extremely short-lived link.

This document does not RECOMMEND sending NFC packets over the Internet or any unsecured network. Especially, there can be a threat model in the scenario of Section 5.1. when the NFC-enabled device links to a NFC-enabled gateway for connectivity with the Internet, the gateway can be attacked. Even though IPv6 over NFC guarantees security between the two NFC devices, there can be another threat during packet forwarding.

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, and Gabriel Montenegro have provided valuable feedback for this draft.

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>>.

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

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

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

Abstract

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

Status of This Memo

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

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

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

This Internet-Draft will expire on September 12, 2019.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

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

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

1. Introduction

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

As IoT (Internet of Things) services become more popular, IPv6 over various link layer technologies such as Bluetooth Low Energy (Bluetooth LE), ITU-T G.9959 (Z-Wave), Digital Enhanced Cordless Telecommunications - Ultra Low Energy (DECT-ULE), Master-Slave/Token Passing (MS/TP), Near Field Communication (NFC), Power Line Communication (PLC), and IEEE 802.15.4e (TSCH), have been defined at [IETF_6lo] working group. IPv6 stacks for constrained node networks use a variation of the 6LoWPAN stack applied to each particular link layer technology.

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

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

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

2. Conventions and Terminology

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

3. 6lo Link layer technologies and possible candidates

3.1. ITU-T G.9959 (specified)

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

3.2. Bluetooth LE (specified)

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

Many Devices such as mobile phones, notebooks, tablets and other handheld computing devices which support Bluetooth 4.0 or subsequent chipsets also support the low-energy variant of Bluetooth. Bluetooth LE is also being included in many different types of accessories that collaborate with mobile devices such as phones, tablets and notebook computers. An example of a use case for a Bluetooth LE accessory is

a heart rate monitor that sends data via the mobile phone to a server on the Internet [RFC7668]. A typical usage of Bluetooth LE is smartphone-based interaction with constrained devices. Bluetooth LE was originally designed to enable star topology networks. However, recent Bluetooth versions support the formation of extended topologies, and IPv6 support for mesh networks of Bluetooth LE devices is being developed [I-D.ietf-6lo-blemesh]

3.3. DECT-ULE (specified)

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

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

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

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

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

3.4. MS/TP (specified)

Master-Slave/Token-Passing (MS/TP) is a Medium Access Control (MAC) protocol for the RS-485 [TIA-485-A] physical layer and is used primarily in building automation networks.

An MS/TP device is typically based on a low-cost microcontroller with limited processing power and memory. These constraints, together

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

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

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

3.5. NFC (specified)

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

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

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

3.6. PLC (specified)

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

The below table shows some available open standards defining PLC.

PLC Systems	Frequency Range	Type	Data Rate	Distance
IEEE1901	<100MHz	Broadband	200Mbps	1000m
IEEE1901.1	<15MHz	PLC-IoT	10Mbps	2000m
IEEE1901.2	<500kHz	Narrowband	200Kbps	3000m

Table 1: Some Available Open Standards in PLC

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

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

3.7. IEEE 802.15.4e (specified)

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

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

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

The main benefits of IEEE 802.15.4 TSCH are:

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

IEEE 802.15.4e can be used for industrial automation.

3.8. Comparison between 6lo Link layer technologies

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

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

Table 2: Comparison between 6lo Link layer technologies

4. 6lo Deployment Scenarios

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

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

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

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

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

jupiterMesh specification is based on the following technologies:

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

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

4.2. Wi-SUN usage of 6lo stacks

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

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

- o Advanced Metering Infrastructure (AMI)
- o Distribution Automation
- o Home Energy Management
- o Infrastructure Management
- o Intelligent Transportation Systems
- o Smart Street Lighting
- o Agriculture
- o Structural health (bridges, buildings etc)
- o Monitoring and Asset Management
- o Smart Thermostats, Air Conditioning and Heat Controls
- o Energy Usage Information Displays

The Wi-SUN Alliance Field Area Network (FAN) covers primarily outdoor networks, and its specification is oriented towards meeting the more rigorous challenges of these environments. Examples include from meter to outdoor access point/router for AMI and DR, or between switches for DA. However, nothing in the profile restricts it to outdoor use. It has the following features;

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

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

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

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

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

- o Smart Metering

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

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

4.4. Netricity usage of 6lo in network layer

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

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

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

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

Netricity system architecture is based on the PHY and MAC layers of IEEE 1901.2 PLC standard. Regarding the 6lo adaptation layer and IPv6 network layer, Netricity utilizes IPv6 protocol suite including 6lo/6LoWPAN header compression, DHCPv6 for IP address management, RPL routing protocol, ICMPv6, and unicast/multicast forwarding. Note that the layer 3 routing in Netricity uses RPL in non-storing mode with the MRHOF objective function based on the own defined Estimated Transmission Time (ETT) metric.

5. Design Space and Guidelines for 6lo Deployment

5.1. Design Space Dimensions for 6lo Deployment

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

- o Deployment/Bootstrapping: 6lo nodes can be connected randomly, or in an organized manner. The bootstrapping has different characteristics for each link layer technology.
- o Topology: Topology of 6lo networks may inherently follow the characteristics of each link layer technology. Point-to-point, star, tree or mesh topologies can be configured, depending on the link layer technology considered.
- o L2-Mesh or L3-Mesh: L2-mesh and L3-mesh may inherently follow the characteristics of each link layer technology. Some link layer technologies may support L2-mesh and some may not support.
- o Multi-link subnet, single subnet: The selection of multi-link subnet and single subnet depends on connectivity and the number of 6lo nodes.

- o Data rate: Typically, the link layer technologies of 6lo have low rate of data transmission. But, by adjusting the MTU, it can deliver higher upper layer data rate.
- o Buffering requirements: Some 6lo use case may require more data rate than the link layer technology support. In this case, a buffering mechanism to manage the data is required.
- o Security and Privacy Requirements: Some 6lo use case can involve transferring some important and personal data between 6lo nodes. In this case, high-level security support is required.
- o Mobility across 6lo networks and subnets: The movement of 6lo nodes depends on the 6lo use case. If the 6lo nodes can move or moved around, a mobility management mechanism is required.
- o Time synchronization requirements: The requirement of time synchronization of the upper layer service is dependent on the 6lo use case. For some 6lo use case related to health service, the measured data must be recorded with exact time and must be transferred with time synchronization.
- o Reliability and QoS: Some 6lo use case requires high reliability, for example real-time service or health-related services.
- o Traffic patterns: 6lo use cases may involve various traffic patterns. For example, some 6lo use case may require short data length and random transmission. Some 6lo use case may require continuous data and periodic data transmission.
- o Security Bootstrapping: Without the external operations, 6lo nodes must have the security bootstrapping mechanism.
- o Power use strategy: to enable certain use cases, there may be requirements on the class of energy availability and the strategy followed for using power for communication [RFC7228]. Each link layer technology defines a particular power use strategy which may be tuned [RFC8352]. Readers are expected to be familiar with [RFC7228] terminology.
- o Update firmware requirements: Most 6lo use cases will need a mechanism for updating firmware. In these cases support for over the air updates are required, probably in a broadcast mode when bandwidth is low and the number of identical devices is high.
- o Wired vs. Wireless: Plenty of 6lo link layer technologies are wireless, except MS/TP and PLC. The selection of wired or wireless link layer technology is mainly dependent on the

requirement of 6lo use cases and the characteristics of wired/wireless technologies. For example, some 6lo use cases may require easy and quick deployment, whereas others may need a continuous source of power.

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

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

- o **Addressing Model:** Addressing model determines whether the device is capable of forming IPv6 Link-local and global addresses and what is the best way to derive the IPv6 addresses for the constrained L2 devices. Whether the device is capable of forming IPv6 Link-local and global addresses, L2-address-derived IPv6 addresses are specified in [RFC4944], but there exist implications for privacy. For global usage, a unique IPv6 address must be derived using an assigned prefix and a unique interface ID. [RFC8065] provides such guidelines. For MAC derived IPv6 address, please refer to [RFC8163] for IPv6 address mapping examples. Broadcast and multicast support are dependent on the L2 networks. Most low-power L2 implementations map multicast to broadcast networks. So care must be taken in the design when to use broadcast and try to stick to unicast messaging whenever possible.
- o **MTU Considerations:** The deployment SHOULD consider their need for maximum transmission unit (MTU) of a packet over the link layer and should consider if fragmentation and reassembly of packets are needed at the 6LoWPAN layer. For example, if the link layer supports fragmentation and reassembly of packets, then 6LoWPAN layer may skip supporting fragmentation/reassembly. In fact, for most efficiency, choosing a low-power link layer that can carry unfragmented application packets would be optimum for packet transmission if the deployment can afford it. Please refer to 6lo RFCs [RFC7668], [RFC8163], [RFC8105] for example guidance.
- o **Mesh or L3-Routing:** 6LoWPAN specifications do provide mechanisms to support for mesh routing at L2. [RFC6550] defines layer three (L3) routing for low power lossy networks using directed graphs. 6LoWPAN is routing protocol agnostic and other L2 or L3 routing protocols can be run using a 6LoWPAN stack.
- o **Address Assignment:** 6LoWPAN requires that IPv6 Neighbor Discovery for low power networks [RFC6775] be used for autoconfiguration of stateless IPv6 address assignment. Considering the energy sensitive networks [RFC6775] makes optimization from classical

IPv6 ND [RFC4861] protocol. It is the responsibility of the deployment to ensure unique global IPv6 addresses for the Internet connectivity. For local-only connectivity IPv6 ULA may be used. [RFC6775] specifies the 6LoWPAN border router(6LBR) which is responsible for prefix assignment to the 6lo/6LoWPAN network. 6LBR can be connected to the Internet or Enterprise network via its one of the interfaces. Please refer to [RFC7668] and [RFC8105] for examples of address assignment considerations. In addition, privacy considerations [RFC8065] must be consulted for applicability. In certain scenarios, the deployment may not support autoconfiguration of IPv6 addressing due to regulatory and business reasons and may choose to offer a separate address assignment service.

- o Header Compression: IPv6 header compression [RFC6282] is a vital part of IPv6 over low power communication. Examples of header compression for different link-layers specifications are found in [RFC7668], [RFC8163], [RFC8105]. A generic header compression technique is specified in [RFC7400].
- o Security and Encryption: Though 6LoWPAN basic specifications do not address security at the network layer, the assumption is that L2 security must be present. In addition, application level security is highly desirable. The working groups [ace] and [core] should be consulted for application and transport level security. 6lo working group is working on address authentication [6lo-ap-nd] and secure bootstrapping is also being discussed at IETF. However, there may be different levels of security available in a deployment through other standards such as hardware level security or certificates for initial booting process. Encryption is important if the implementation can afford it.
- o Additional processing: [RFC8066] defines guidelines for ESC dispatch octets use in the 6LoWPAN header. An implementation may take advantage of ESC header to offer a deployment specific processing of 6LoWPAN packets.

6. 6lo Use Case Examples

As IPv6 stacks for constrained node networks use a variation of the 6LoWPAN stack applied to each particular link layer technology, various 6lo use cases can be provided. In this clause, one 6lo use case example of Bluetooth LE (Smartphone-Based Interaction with Constrained Devices) is described. Other 6lo use case examples are described in Appendix.

The key feature behind the current high Bluetooth LE momentum is its support in a large majority of smartphones in the market. Bluetooth

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

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

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

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

7. IANA Considerations

There are no IANA considerations related to this document.

8. Security Considerations

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

9. Acknowledgements

Carles Gomez has been funded in part by the Spanish Government through the Jose Castillejo CAS15/00336 grant, and through the TEC2016-79988-P grant. His contribution to this work has been carried out in part during his stay as a visiting scholar at the Computer Laboratory of the University of Cambridge.

Thomas Watteyne, Pascal Thubert, Xavier Vilajosana, Daniel Migault, and Jianqiang HOU have provided valuable feedback for this draft.

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

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

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4919] Kushalnagar, N., Montenegro, G., and C. Schumacher, "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals", RFC 4919, DOI 10.17487/RFC4919, August 2007, <<https://www.rfc-editor.org/info/rfc4919>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<https://www.rfc-editor.org/info/rfc4944>>.
- [RFC5826] Brandt, A., Buron, J., and G. Porcu, "Home Automation Routing Requirements in Low-Power and Lossy Networks", RFC 5826, DOI 10.17487/RFC5826, April 2010, <<https://www.rfc-editor.org/info/rfc5826>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<https://www.rfc-editor.org/info/rfc6282>>.

- [RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <<https://www.rfc-editor.org/info/rfc6550>>.
- [RFC6568] Kim, E., Kaspar, D., and JP. Vasseur, "Design and Application Spaces for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6568, DOI 10.17487/RFC6568, April 2012, <<https://www.rfc-editor.org/info/rfc6568>>.
- [RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <<https://www.rfc-editor.org/info/rfc6775>>.
- [RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <<https://www.rfc-editor.org/info/rfc7228>>.
- [RFC7400] Bormann, C., "6LoWPAN-GHC: Generic Header Compression for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 7400, DOI 10.17487/RFC7400, November 2014, <<https://www.rfc-editor.org/info/rfc7400>>.
- [RFC7428] Brandt, A. and J. Buron, "Transmission of IPv6 Packets over ITU-T G.9959 Networks", RFC 7428, DOI 10.17487/RFC7428, February 2015, <<https://www.rfc-editor.org/info/rfc7428>>.
- [RFC7554] Watteyne, T., Ed., Palattella, M., and L. Grieco, "Using IEEE 802.15.4e Time-Slotted Channel Hopping (TSCH) in the Internet of Things (IoT): Problem Statement", RFC 7554, DOI 10.17487/RFC7554, May 2015, <<https://www.rfc-editor.org/info/rfc7554>>.
- [RFC7668] Nieminen, J., Savolainen, T., Isomaki, M., Patil, B., Shelby, Z., and C. Gomez, "IPv6 over BLUETOOTH(R) Low Energy", RFC 7668, DOI 10.17487/RFC7668, October 2015, <<https://www.rfc-editor.org/info/rfc7668>>.

- [RFC8036] Cam-Winget, N., Ed., Hui, J., and D. Popa, "Applicability Statement for the Routing Protocol for Low-Power and Lossy Networks (RPL) in Advanced Metering Infrastructure (AMI) Networks", RFC 8036, DOI 10.17487/RFC8036, January 2017, <<https://www.rfc-editor.org/info/rfc8036>>.
- [RFC8065] Thaler, D., "Privacy Considerations for IPv6 Adaptation-Layer Mechanisms", RFC 8065, DOI 10.17487/RFC8065, February 2017, <<https://www.rfc-editor.org/info/rfc8065>>.
- [RFC8066] Chakrabarti, S., Montenegro, G., Droms, R., and J. Woodyatt, "IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) ESC Dispatch Code Points and Guidelines", RFC 8066, DOI 10.17487/RFC8066, February 2017, <<https://www.rfc-editor.org/info/rfc8066>>.
- [RFC8105] Mariager, P., Petersen, J., Ed., Shelby, Z., Van de Logt, M., and D. Barthel, "Transmission of IPv6 Packets over Digital Enhanced Cordless Telecommunications (DECT) Ultra Low Energy (ULE)", RFC 8105, DOI 10.17487/RFC8105, May 2017, <<https://www.rfc-editor.org/info/rfc8105>>.
- [RFC8163] Lynn, K., Ed., Martocci, J., Neilson, C., and S. Donaldson, "Transmission of IPv6 over Master-Slave/Token-Passing (MS/TP) Networks", RFC 8163, DOI 10.17487/RFC8163, May 2017, <<https://www.rfc-editor.org/info/rfc8163>>.
- [RFC8352] Gomez, C., Kovatsch, M., Tian, H., and Z. Cao, Ed., "Energy-Efficient Features of Internet of Things Protocols", RFC 8352, DOI 10.17487/RFC8352, April 2018, <<https://www.rfc-editor.org/info/rfc8352>>.

10.2. Informative References

- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <<https://www.rfc-editor.org/info/rfc3315>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.

- [I-D.ietf-6lo-nfc]
Choi, Y., Hong, Y., Youn, J., Kim, D., and J. Choi,
"Transmission of IPv6 Packets over Near Field
Communication", draft-ietf-6lo-nfc-13 (work in progress),
February 2019.
- [I-D.ietf-roll-aodv-rpl]
Anamalamudi, S., Zhang, M., Perkins, C., Anand, S., and B.
Liu, "Asymmetric AODV-P2P-RPL in Low-Power and Lossy
Networks (LLNs)", draft-ietf-roll-aodv-rpl-06 (work in
progress), March 2019.
- [I-D.ietf-6tisch-6top-sfx]
Dujovne, D., Grieco, L., Palattella, M., and N. Accettura,
"6TiSCH Experimental Scheduling Function (SFX)", draft-
ietf-6tisch-6top-sfx-01 (work in progress), March 2018.
- [I-D.ietf-6lo-blemesh]
Gomez, C., Darroudi, S., Savolainen, T., and M. Spoerk,
"IPv6 Mesh over BLUETOOTH(R) Low Energy using IPSP",
draft-ietf-6lo-blemesh-04 (work in progress), January
2019.
- [I-D.satish-6tisch-6top-sf1]
Anamalamudi, S., Liu, B., Zhang, M., Sangi, A., Perkins,
C., and S. Anand, "Scheduling Function One (SF1): hop-by-
hop Scheduling with RSVP-TE in 6tisch Networks", draft-
satish-6tisch-6top-sf1-04 (work in progress), October
2017.
- [I-D.ietf-6lo-plc]
Hou, J., Liu, B., Hong, Y., Tang, X., and C. Perkins,
"Transmission of IPv6 Packets over PLC Networks", draft-
ietf-6lo-plc-00 (work in progress), February 2019.
- [IETF_6lo]
"IETF IPv6 over Networks of Resource-constrained Nodes
(6lo) working group",
<<https://datatracker.ietf.org/wg/6lo/charter/>>.
- [TIA-485-A]
"TIA, "Electrical Characteristics of Generators and
Receivers for Use in Balanced Digital Multipoint Systems",
TIA-485-A (Revision of TIA-485)", March 2003,
<[https://global.ihs.com/
doc_detail.cfm?item_s_key=00032964](https://global.ihs.com/doc_detail.cfm?item_s_key=00032964)>.
- [G3-PLC] "G3-PLC Alliance", <<http://www.g3-plc.com/home/>>.

- [NETRICITY] "Netricity program in HomePlug Powerline Alliance", <<http://groups.homeplug.org/tech/Netricity>>.
- [G.9959] "International Telecommunication Union, "Short range narrow-band digital radiocommunication transceivers - PHY and MAC layer specifications", ITU-T Recommendation", January 2015.
- [G.9903] "International Telecommunication Union, "Narrowband orthogonal frequency division multiplexing power line communication transceivers for G3-PLC networks", ITU-T Recommendation", August 2017.
- [IEEE1901] "IEEE Standard, IEEE Std. 1901-2010 - IEEE Standard for Broadband over Power Line Networks: Medium Access Control and Physical Layer Specifications", 2010, <<https://standards.ieee.org/findstds/standard/1901-2010.html>>.
- [IEEE1901.1] "IEEE Standard (work-in-progress), IEEE-SA Standards Board", <<http://sites.ieee.org/sagroups-1901-1/>>.
- [IEEE1901.2] "IEEE Standard, IEEE Std. 1901.2-2013 - IEEE Standard for Low-Frequency (less than 500 kHz) Narrowband Power Line Communications for Smart Grid Applications", 2013, <<https://standards.ieee.org/findstds/standard/1901.2-2013.html>>.
- [BACnet] "ASHRAE, "BACnet-A Data Communication Protocol for Building Automation and Control Networks", ANSI/ASHRAE Standard 135-2016", January 2016, <http://www.techstreet.com/ashrae/standards/ashrae-135-2016?product_id=1918140#jumps>.

Appendix A. Other 6lo Use Case Examples

A.1. Use case of ITU-T G.9959: Smart Home

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

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

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

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

A.2. Use case of DECT-ULE: Smart Home

DECT is a technology widely used for wireless telephone communications in residential scenarios. Since DECT-ULE is a low-power variant of DECT, DECT-ULE can be used to connect constrained devices such as sensors and actuators to a Fixed Part, a device that typically acts as a base station for wireless telephones. Therefore, DECT-ULE is specially suitable for the connected home space in application areas such as home automation, smart metering, safety, healthcare, etc.

Example: Use of DECT-ULE for Smart Metering

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

A.3. Use case of MS/TP: Building Automation Networks

The primary use case for IPv6 over MS/TP (6LoBAC) is in building automation networks. [BACnet] is the open international standard protocol for building automation, and MS/TP is defined in [BACnet] Clause 9. MS/TP was designed to be a low cost multi-drop field bus to inter-connect the most numerous elements (sensors and actuators) of a building automation network to their controllers. A key aspect of 6LoBAC is that it is designed to co-exist with BACnet MS/TP on the same link, easing the ultimate transition of some BACnet networks to native end-to-end IPv6 transport protocols. New applications for 6LoBAC may be found in other domains where low cost, long distance, and low latency are required.

Example: Use of 6LoBAC in Building Automation Networks

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

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

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

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

According to applications, various secured data can be handled and transferred. Depending on security level of the data, methods for transfer can be alternatively selected.

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

A senior citizen who lives alone wears one to several wearable 6lo devices to measure heartbeat, pulse rate, etc. The 6lo devices are densely installed at home for movement detection. An LoWPAN Border Router (LBR) at home will send the sensed information to a connected

healthcare center. Portable base stations with LCDs may be used to check the data at home, as well. Data is gathered in both periodic and event-driven fashion. In this application, event-driven data can be very time-critical. In addition, privacy also becomes a serious issue in this case, as the sensed data is very personal.

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

A.5. Use case of PLC: Smart Grid

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

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

Example: Use of PLC for Advanced Metering Infrastructure

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

With the existing power line infrastructure as communication medium, cost on building up the PLC network is naturally saved, and more importantly, labor operational costs can be minimized from a long-

term perspective. Furthermore, this AMI application speeds up electricity charge, reduces losses by restraining power theft and helps to manage the health of the grid based on line loss analysis.

Example: Use of PLC (IEEE1901.1) for WASA in Smart Grid

Many sub-systems of Smart Grid require low data rate and narrowband variant (IEEE1901.2) of PLC fulfills such requirements. Recently, more complex scenarios are emerging that require higher data rates.

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

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

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

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

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

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

Authors' Addresses

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

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

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

Email: carlesgo@entel.upc.edu

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

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

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

Email: sangi_bahrian@yahoo.com

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

Email: take@modio.se

Samita Chakrabarti
San Jose, CA
USA

Email: samitac.ietf@gmail.com

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

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

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

Abstract

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

Status of This Memo

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

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

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

This Internet-Draft will expire on July 29, 2019.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1. Introduction	2
2. Terminology	3
2.1. BCP 14	3
2.2. References	4
2.3. New Terms	4
2.4. Acronym Definitions	4
3. Overview	5
4. Updating RFC 8505	7
4.1. Extended Neighbor Discovery Options and Messages	7
4.1.1. Extending the Capability Indication Option	8
4.1.2. New Code Prefix for Address Mapping Messages	8
4.1.3. New ARO Status	8
4.2. Address Mapping Messages	9
4.3. IPv6 ND-based Address Lookup	10
5. Backward Compatibility	10
6. Security Considerations	10
7. IANA Considerations	10
7.1. ICMP Codes	11
7.2. New ARO Status values	11
7.3. New 6LoWPAN Capability Bits	12
8. Acknowledgments	12
9. References	12
9.1. Normative References	12
9.2. Informative References	13
Authors' Addresses	14

1. Introduction

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

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

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

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

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

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

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

2. Terminology

2.1. BCP 14

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

2.2. References

This document uses terms and concepts that are discussed in:

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

2.3. New Terms

This document introduces the following terminology:

Address Mapping Request

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

Address Mapping Confirm

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

Address Registrar

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

Address Registration

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

2.4. Acronym Definitions

This document uses the following acronyms:

6BBR: 6LoWPAN Backbone Router

6LBR: 6LoWPAN Border Router

6LR: 6LoWPAN Router

6CIO: Capability Indication Option

AMC: Address Mapping Confirmation

AMR: Address Mapping Request

ARO: Address Registration Option

DAC: Duplicate Address Confirmation

DAD: Duplicate Address Detection

DAR: Duplicate Address Request

EDAC: Extended Duplicate Address Confirmation

EDAR: Extended Duplicate Address Request

DODAG: Destination-Oriented Directed Acyclic Graph

LLN: Low-Power and Lossy Network

NA: Neighbor Advertisement

NCE: Neighbor Cache Entry

ND: Neighbor Discovery

NS: Neighbor Solicitation

ROVR: Registration Ownership Verifier

RA: Router Advertisement

RS: Router Solicitation

TID: Transaction ID

3. Overview

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

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

The LLN may be a hub-and-spoke access link such as (Low-Power) IEEE STD. 802.11 (Wi-Fi) [IEEEstd80211] and IEEE STD. 802.15.1 (Bluetooth) [IEEEstd802151], or a Mesh-Under or a Route-Over network [RFC8505].




Figure 1: Backbone Link and 6LBR

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

The Address Registration retains:

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

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

Unless otherwise configured, a 6LBR does the following:

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

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

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

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

4. Updating RFC 8505

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

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

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

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

4.1. Extended Neighbor Discovery Options and Messages

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

4.1.1.1. Extending the Capability Indication Option

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

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




Figure 2: New Capability Bits in the 6CIO

Option Fields:

Type: 36

A: The 6LBR provides address mapping services.

4.1.1.2. New Code Prefix for Address Mapping Messages

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

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

4.1.1.3. New ARO Status

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

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

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

Table 1: EARO Status

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

4.2. Address Mapping Messages

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

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

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

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

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

4.3. IPv6 ND-based Address Lookup

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

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

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

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

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

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

5. Backward Compatibility

6. Security Considerations

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

7. IANA Considerations

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

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

7.1. ICMP Codes

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

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

The new subregistries are to be initialized as follows:

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

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

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

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

7.2. New ARO Status values

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

ARO Status	Description	Document
11	Not Found	This RFC

Table 4: New ARO Status values

7.3. New 6LoWPAN Capability Bits

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

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

Table 5: New 6LoWPAN Capability Bits

8. Acknowledgments

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <<https://www.rfc-editor.org/info/rfc4862>>.
- [RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <<https://www.rfc-editor.org/info/rfc6775>>.

- [RFC7400] Bormann, C., "6LoWPAN-GHC: Generic Header Compression for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 7400, DOI 10.17487/RFC7400, November 2014, <<https://www.rfc-editor.org/info/rfc7400>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8505] Thubert, P., Ed., Nordmark, E., Chakrabarti, S., and C. Perkins, "Registration Extensions for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Neighbor Discovery", RFC 8505, DOI 10.17487/RFC8505, November 2018, <<https://www.rfc-editor.org/info/rfc8505>>.

9.2. Informative References

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

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

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

Authors' Addresses

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

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

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

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

6MAN
Internet-Draft
Intended status: Informational
Expires: November 3, 2019

P. Thubert, Ed.
Cisco Systems
May 2, 2019

IPv6 Neighbor Discovery on Wireless Networks
draft-thubert-6man-ipv6-over-wireless-03

Abstract

This document describes how the original IPv6 Neighbor Discovery and Wireless ND (WiND) can be applied on various abstractions of wireless media.

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 November 3, 2019.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1. Introduction	2
2. Acronyms	4
3. IP Models	6
3.1. Physical Broadcast Domain	6
3.2. MAC-Layer Broadcast Emulations	7
3.3. Mapping the IPv6 Link Abstraction	8
3.4. Mapping the IPv6 Subnet Abstraction	9
4. Wireless ND	10
4.1. Introduction to WiND	10
4.2. Links and Link-Local Addresses	11
4.3. Subnets and Global Addresses	11
5. WiND Applicability	12
5.1. Case of LPWANs	13
5.2. Case of Infrastructure BSS and ESS	13
5.3. Case of Mesh Under Technologies	14
5.4. Case of DMC radios	14
5.4.1. Using IPv6 ND only	15
5.4.2. Using Wireless ND	15
6. IANA Considerations	17
7. Security Considerations	18
8. Acknowledgments	18
9. References	18
9.1. Normative References	18
9.2. Informative References	19
Author's Address	21

1. Introduction

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

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

Wi-Fi [IEEEstd80211] Access Points (APs) deployed in an Extended Service Set (ESS) act as Ethernet Bridges [IEEEstd8021], with the property that the bridging state is established at the time of association. This ensures connectivity to the node (STA) and protects the wireless medium against broadcast-intensive Transparent Bridging reactive Lookups.

In other words, the association process is used to register the MAC Address of the STA to the AP. The AP subsequently proxies the bridging operation and does not need to forward the broadcast Lookups over the radio.

Like Transparent Bridging, IPv6 [RFC8200] Neighbor Discovery [RFC4861] [RFC4862] Protocol (IPv6 ND) is a reactive protocol, based on multicast transmissions to locate an on-link correspondent and ensure the uniqueness of an IPv6 address. The mechanism for Duplicate Address Detection (DAD) [RFC4862] was designed for the efficient broadcast operation of Ethernet Bridging. Since broadcast can be unreliable over wireless media, DAD often fails to discover duplications [I-D.yourtchenko-6man-dad-issues]. In practice, IPv6 addresses very rarely conflict because of the entropy of the 64-bit Interface IDs, not because address duplications are detected and resolved.

The IPv6 ND Neighbor Solicitation (NS) [RFC4861] message is used for DAD and address Lookup when a node moves, or wakes up and reconnects to the wireless network. The NS message is targeted to a Solicited-Node Multicast Address (SNMA) [RFC4291] and should in theory only reach a very small group of nodes. But in reality, IPv6 multicast messages are typically broadcast on the wireless medium, and so they are processed by most of the wireless nodes over the subnet (e.g., the ESS fabric) regardless of how few of the nodes are subscribed to the SNMA. As a result, IPv6 ND address Lookups and DADs over a large wireless and/or a LowPower Lossy Network (LLN) can consume enough bandwidth to cause a substantial degradation to the unicast traffic service [I-D.vyncke-6man-mcast-not-efficient].

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

These problems can be alleviated by reducing the IPv6 ND broadcasts over wireless access links. This has been done by splitting the broadcast domains and by routing between subnets, at the extreme by assigning a /64 prefix to each wireless node (see [RFC8273]).

Another way is to proxy at the boundary of the wired and wireless domains the Layer-3 protocols that rely on MAC Layer broadcast operations. For instance, IEEE 802.11 [IEEEstd80211] situates proxy-ARP (IPv4) and proxy-ND (IPv6) functions at the Access Points (APs).

But proxying ND requires a perfect knowledge of the peer IPv6 addresses for which proxying is provided. In a generic fashion, radio connectivity changes with movements and variations in the environment, which makes forming and maintaining that knowledge a hard problem in the general case.

Discovering peer addresses by snooping the IPV6 ND protocol as proposed for SAVI [I-D.bi-savi-wlan] was found to be unreliable. An IPv6 address may not be discovered immediately due to a packet loss, or if a "silent" node is not currently using one of its addresses, e.g., a node that waits in wake-on-lan state. A change of state, e.g. due to a movement, may be missed or misordered, leading to unreliable connectivity and an incomplete knowledge of the set of peers.

Wireless ND (WiND) introduces a new approach to IPv6 ND that is designed to apply to the WLANs and WPANs types of networks. On the one hand, WiND avoids the use of broadcast operation for Address Resolution and Duplicate Address Detection, and on the other hand, WiND supports use cases where Subnet and MAC-level domains are not congruent, which is common in those types of networks unless a specific MAC-Level emulation is provided.

To achieve this, WiND applies routing inside the Subnets, which enables MultiLink Subnets. Hosts register their addresses to their serving routers with [RFC8505]. With the registration, routers have a complete knowledge of the hosts they serve and in return, hosts obtain routing services for their registered addresses. The registration is abstract to the routing protocol, and it can be protected to prevent impersonation attacks with [I-D.ietf-6lo-ap-nd].

The routing service can be a simple reflexion in a Hub-and-Spoke Subnet that emulates an IEEE Std 802.11 Infrastructure BSS at Layer 3. It can also be a full-fledge routing protocol, in particular RPL [RFC6550] that was designed to adapt to various LLNs such as WLAN and WPAN radio meshes with the concept of Objective Function. Finally, the routing service can also be ND proxy that emulates an IEEE Std 802.11 Infrastructure ESS at Layer 3. WiND specifies the IPv6 Backbone Router for that purpose in [I-D.ietf-6lo-backbone-router].

More details on WiND can be found in Section 4.1.

2. Acronyms

This document uses the following abbreviations:

6BBR: 6LoWPAN Backbone Router

6LBR: 6LoWPAN Border Router

6LN: 6LoWPAN Node

6LR: 6LoWPAN Router

ARO: Address Registration Option

DAC: Duplicate Address Confirmation

DAD: Duplicate Address Detection

DAR: Duplicate Address Request

EDAC: Extended Duplicate Address Confirmation

EDAR: Extended Duplicate Address Request

MLSN: Multi-Link Subnet

LLN: Low-Power and Lossy Network

NA: Neighbor Advertisement

NBMA: Non-Broadcast Multi-Access

NCE: Neighbor Cache Entry

ND: Neighbor Discovery

NDP: Neighbor Discovery Protocol

NS: Neighbor Solicitation

RPL: IPv6 Routing Protocol for LLNs

RA: Router Advertisement

RS: Router Solicitation

WiND: Wireless Neighbor Discovery

WLAN: Wireless Local Area Network

WPAN: Wireless Personal Area Network

3. IP Models

3.1. Physical Broadcast Domain

At the physical (PHY) Layer, a broadcast domain is the set of nodes that may receive a datagram that one sends over an interface, in other words the set of nodes in range of radio transmission. This set can comprise a single peer on a serial cable used as point-to-point (P2P) link. It may also comprise multiple peer nodes on a broadcast radio or a shared physical resource such as the legacy Ethernet shared wire.

On WLAN and WPAN radios, the physical broadcast domain is defined by a particular transmitter, as the set of nodes that can receive what this transmitter is sending. Literally every datagram defines its own broadcast domain since the chances of reception of a given datagram are statistical. In average and in stable conditions, the broadcast domain of a particular node can be still be seen as mostly constant and can be used to define a closure of nodes on which an upper-layer abstraction can be built.

A PHY-layer communication can be established between 2 nodes if their physical broadcast domains overlap.

On WLAN and WPAN radios, this property is usually reflexive, meaning that if B can receive a datagram from A, then A can receive a datagram from B. But there can be asymmetries due to power levels, interferers near one of the receivers, or differences in the quality of the hardware (e.g., crystals, PAs and antennas) that may affect the balance to the point that the connectivity becomes mostly unidirectional, e.g., A to B but practically not B to A. It takes a particular effort to place a set of devices in a fashion that all their physical broadcast domains fully overlap, and it can not be assumed in the general case. In other words, the property of radio connectivity is generally not transitive, meaning that A may be in range with B and B may be in range with C does not necessarily imply that A is in range with C.

We define MAC-Layer Direct Broadcast (DMC) a transmission mode where the broadcast domain that is usable at the MAC layer is directly the physical broadcast domain. IEEE 802.15.4 [IEEE802154] and IEEE 802.11 [IEEEstd80211] OCB (for Out of the Context of a BSS) are examples of DMC radios. This contrasts with a number of MAC-layer Broadcast Emulation schemes that are described in the next section.

3.2. MAC-Layer Broadcast Emulations

While a physical broadcast domain is constrained to a single shared wire, Ethernet Bridging emulates the broadcast properties of that wire over a whole physical mesh of Ethernet links. For the upper layer, the qualities of the shared wire are essentially conserved, with a reliable and cheap broadcast operation over a closure of nodes defined by their connectivity to the emulated wire.

In large switched fabrics, overlay techniques enable a limited connectivity between nodes that are known to a mapping server. The emulated broadcast domain is configured to the system, e.g., with a VXLAN network identifier (VNID). Broadcast operations on the overlay can be emulated but can become very expensive, and it makes sense to proactively install the relevant state in the mapping server as opposed to rely on reactive broadcast lookups.

An IEEE Std 802.11 Infrastructure Basic Service Set (BSS) also provides a closure of nodes as defined by the broadcast domain of a central Access Point (AP). The AP relays both unicast and broadcast packets and ensures a reflexive and transitive emulation of the shared wire between the associated nodes, with the capability to signal link-up/link-down to the upper layer. Within an Infrastructure BSS, the physical broadcast domain of the AP serves as emulated broadcast domain for all the nodes that are associated to the AP. Broadcast packets are relayed by the AP and are not acknowledged. To ensure that all nodes in the BSS receive the broadcast transmission, AP transmits at the slowest PHY speed. This translates into maximum co-channel interferences for others and longest occupancy of the medium, for a duration that can be 100 times that of a unicast. For that reason, upper layer protocols should tend to avoid the use of broadcast when operating over Wi-Fi.

In an IEEE Std 802.11 Infrastructure Extended Service Set (ESS), infrastructure BSSes are interconnected by a bridged network, typically running Transparent Bridging and Spanning tree Protocol. In the original model, the state in the Transparent Bridge is set by observing the source MAC address of the frames. When a state is missing for a destination MAC address, the frame is broadcasted with the expectation that the response will populate the state. This is a reactive operation, meaning that the state is populated reactively to a need for forwarding. It is also possible to send a gratuitous frame to advertise self throughout the bridged network, and that is also a broadcast. The process of the association prepares a bridging state proactively at the AP, so as to avoid the reactive broadcast lookup. It may also generate a gratuitous broadcast sourced at the MAC address of the STA to prepare or update the state in the Transparent Bridges. This model avoids the need of multicast over

the wireless access, and it is only logical that IPv6 ND evolved towards proposes similar methods at Layer-3 for its operation.

In some cases of WLAN and WPAN radios, a mesh-under technology (e.g., a IEEE 802.11s or IEEE 802.15.10) provides meshing services that are similar to bridging, and the broadcast domain is well defined by the membership of the mesh. Mesh-Under emulates a broadcast domain by flooding the broadcast packets at Layer-2. When operating on a single frequency, this operation is known to interfere with itself, forcing deployment to introduce delays that dampen the collisions. All in all, the mechanism is slow, inefficient and expensive.

Going down the list of cases above, the cost of a broadcast transmissions becomes increasingly expensive, and there is a push to rethink the upper-layer protocols so as to reduce the dependency on broadcast operations.

There again, a MAC-layer communication can be established between 2 nodes if their MAC-layer broadcast domains overlap. In the absence of a MAC-layer emulation such as a mesh-under or an Infrastructure BSS, the MAC-layer broadcast domain is congruent with that of the PHY-layer and inherits its properties for reflexivity and transitivity. IEEE 802.11p, which operates Out of the Context of a BSS (DMC radios) is an example of a network that does not have a MAC-Layer broadcast domain emulation, which means that it will exhibit mostly reflexive and mostly non-transitive transmission properties.

3.3. Mapping the IPv6 Link Abstraction

IPv6 defines a concept of Link, Link Scope and Link-Local Addresses (LLA), an LLA being unique and usable only within the Scope of a Link. The IPv6 Neighbor Discovery (ND) [RFC4861][RFC4862] Duplicate Address Detection (DAD) process leverages a multicast transmission to ensure that an IPv6 address is unique as long as the owner of the address is connected to the broadcast domain. It must be noted that in all the cases in this specification, the Layer-3 multicast operation is always a MAC_Layer broadcast for the lack of a Layer-2 multicast operation that could handle a possibly very large number of groups in order to make the unicast efficient. This means that for every multicast packet regardless of the destination group, all nodes will receive the packet and process it all the way to Layer-3.

On wired media, the Link is often confused with the physical broadcast domain because both are determined by the serial cable or the Ethernet shared wire. Ethernet Bridging reinforces that illusion by providing a MAC-Layer broadcast domain that emulates a physical broadcast domain over the mesh of wires. But the difference shows on legacy Non-Broadcast Multi-Access (NBMA) such as ATM and Frame-Relay,

on shared links and on newer types of NBMA networks such as radio and composite radio-wires networks. It also shows when private VLANs or Layer-2 cryptography restrict the capability to read a frame to a subset of the connected nodes.

In mesh-under and Infrastructure BSS, the IP Link extends beyond the physical broadcast domain to the emulated MAC-Layer broadcast domain. Relying on Multicast for the ND operation remains feasible but becomes detrimental to unicast traffic, energy-inefficient and unreliable, and its use is discouraged.

On DMC radios, IP Links between peers come and go as the individual physical broadcast domains of the transmitters meet and overlap. The DAD operation cannot provide once and for all guarantees on the broadcast domain defined by one radio transmitter if that transmitter keeps meeting new peers on the go. The nodes may need to form new LLAs to talk to one another and the scope where LLA uniqueness can be dynamically checked is that pair of nodes. As long as there's no conflict a node may use the same LLA with multiple peers but it has to revalidate DAD with every new peer node. In practice, each pair of nodes defines a temporary P2P link, which can be modeled as a sub-interface of the radio interface.

3.4. Mapping the IPv6 Subnet Abstraction

IPv6 also defines a concept of Subnet for Glocal and Unique Local Addresses. Addresses in a same Subnet share a same prefix and by extension, a node belongs to a Subnet if it has an interface with an address on that Subnet. A Subnet prefix is Globally Unique so it is sufficient to validate that an address that is formed from a Subnet prefix is unique within that Subnet to guarantee that it is globally unique. IPv6 aggregation relies on the property that a packet from the outside of a Subnet can be routed to any router that belongs to the Subnet, and that this router will be able to either resolve the destination MAC address and deliver the packet, or route the packet to the destination within the Subnet. If the Subnet is known as onlink, then any node may also resolve the destination MAC address and deliver the packet, but if the Subnet is not onlink, then a host that does not have an NCE for the destination will need to pass the packet to a router.

On IEEE Std. 802.3, a Subnet is often congruent with an IP Link because both are determined by the physical attachment to an Ethernet shared wire or an IEEE Std. 802.1 bridged broadcast domain. In that case, the connectivity over the Link is transitive, the Subnet can appear as onlink, and any node can resolve a destination MAC address of any other node directly using IPv6 Neighbor Discovery.

But an IP Link and an IP Subnet are not always congruent. In a shared Link situation, a Subnet may encompass only a subset of the nodes connected to the Link. In Route-Over Multi-Link Subnets (MLSN) [RFC4903], routers federate the Links between nodes that belong to the Subnet, the Subnet is not onlink and it extends beyond any of the federated Links.

The DAD and lookup procedures in IPv6 ND expects that a node in a Subnet is reachable within the broadcast domain of any other node in the Subnet when that other node attempts to form an address that would be a duplicate or attempts to resolve the MAC address of this node. This is why ND is only applicable for P2P and transit links, and requires extensions for other topologies.

4. Wireless ND

4.1. Introduction to WiND

Wireless Neighbor Discovery (WiND) [RFC6775] [RFC8505] [I-D.ietf-6lo-backbone-router] [I-D.ietf-6lo-ap-nd] defines a new ND operation that is based on 2 major paradigm changes, proactive address registration by hosts to their attachment routers and routing to host routes (/128) within the subnet. This allows WiND to avoid the classical ND expectations of transit links and Subnet-wide broadcast domains.

The proactive address registration is performed with a new option in NS/NA messages, the Extended Address Registration Option (EARO) defined in [RFC8505]. This method allows to prepare and maintain the host routes in the routers and avoids the reactive NS(Lookup) found in IPv6 ND. This is a direct benefit for wireless Links since it avoids the MAC level broadcasts that are associated to NS(Lookup).

The EARO provides information to the router that is independent to the routing protocol and routing can take multiple forms, from a traditional IGP to a collapsed ub-and-Spoke model where only one router owns and advertises the prefix. [RFC8505] is already referenced for RIFT [I-D.ietf-rift-rift], RPL [RFC6550] with [I-D.thubert-roll-unaware-leaves] and IPv6 ND proxy [I-D.ietf-6lo-backbone-router].

WiND does not change IPv6 addressing [RFC4291] or the current practices of assigning prefixes to subnets. It is still typical to assign a /64 to a subnet and to use interface IDs of 64 bits. Duplicate Address detection within the Subnet is performed with a central registrar, using new ND Extended Duplicate Address messages (EDAR and EDAC) [RFC8505]. This operation modernizes ND for application in overlays with Map Resolvers and enables unicast

lookups [I-D.thubert-6lo-unicast-lookup] for addresses registered to the resolver.

WiND also enables to extend a legacy /64 on Ethernet with ND proxy over the wireless. This way nodes can form any address they want and move freely from an L3-AP (that is really a backbone router in bridging mode, more in [I-D.ietf-6lo-backbone-router]) to another, without renumbering. Backbone Routers federate multiple LLNs over a Backbone Link to form a MultiLink Subnet (MLSN). Backbone Routers placed along the LLN edge of the Backbone handle IPv6 Neighbor Discovery, and forward packets on behalf of registered nodes.

An LLN node (6LN) registers all its IPv6 Addresses using an NS(EARO) as specified in [RFC8505] to the 6BBR. The 6BBR is also a Border Router that performs IPv6 Neighbor Discovery (IPv6 ND) operations on its Backbone interface on behalf of the 6LNs that have registered addresses on its LLN interfaces without the need of a broadcast over the wireless medium.

WiND is also compatible with DHCPv6 and other forms of address assignment in which case it can still be used for DAD.

4.2. Links and Link-Local Addresses

For Link-Local Addresses, DAD is performed between communicating pairs of nodes. It is carried out as part of a registration process that is based on a NS/NA exchange that transports an EARO. During that process, the DAD is validated and a Neighbor Cache Entry (NCE) is populated with a single unicast exchange.

For instance, in the case of a Bluetooth Low Energy (BLE) [RFC7668][IEEEstd802151] Hub-and-Spoke configuration, Uniqueness of Link local Addresses need only to be verified between the pairs of communicating nodes, a central router and a peripheral host. In that example, 2 peripheral hosts connected to the same central router can not have the same Link Local Address because the Binding Cache Entries (BCEs) would collide at the central router which could not talk to both over the same interface. The WiND operation is appropriate for that DAD operation, but the one from ND is not, because peripheral hosts are not on the same broadcast domain. On the other hand, Global and ULA DAD is validated at the Subnet Level, using a registrar hosted by the central router.

4.3. Subnets and Global Addresses

WiND extends IPv6 ND for Hub-and-Spoke (e.g., BLE) and Route-Over (e.g., RPL) Multi-Link Subnets (MLSNs).

In the Hub-and-Spoke case, each Hub-Spoke pair is a distinct IP Link, and a Subnet can be mapped on a collection of Links that are connected to the Hub. The Subnet prefix is associated to the Hub. Acting as 6LR, the Hub advertises the prefix as not-onlink to the spokes in RA messages Prefix Information Options (PIO). Acting as 6LNs, the Spokes autoconfigure addresses from that prefix and register them to the Hub with a corresponding lifetime. Acting as a 6LBR, the Hub maintains a binding table of all the registered IP addresses and rejects duplicate registrations, thus ensuring a DAD protection for a registered address even if the registering node is sleeping. Acting as 6LR, the Hub also maintains an NCE for the registered addresses and can deliver a packet to any of them for their respective lifetimes. It can be observed that this design builds a form of Layer-3 Infrastructure BSS.

A Route-Over MLSN is considered as a collection of Hub-and-Spoke where the Hubs form a connected dominating set of the member nodes of the Subnet, and IPv6 routing takes place between the Hubs within the Subnet. A single logical 6LBR is deployed to serve the whole mesh. The registration in [RFC8505] is abstract to the routing protocol and provides enough information to feed a routing protocol such as RPL as specified in [I-D.thubert-roll-unaware-leaves]. In a degraded mode, all the Hubs are connected to a same high speed backbone such as an Ethernet bridging domain where IPv6 ND is operated. In that case, it is possible to federate the Hub, Spoke and Backbone nodes as a single Subnet, operating IPv6 ND proxy operations [I-D.ietf-6lo-backbone-router] at the Hubs, acting as 6BBRs. It can be observed that this latter design builds a form of Layer-3 Infrastructure ESS.

5. WiND Applicability

WiND allows P2P, P2MP hub-and spoke, MAC-level broadcast domain emulation such as mesh-under and Wi-Fi BSS, and Route-Over meshes.

There is an intersection where Link and Subnet are congruent and where both ND and WiND could apply. These includes P2P, the MAC emulation of a PHY broadcast domain, and the particular case of always on, fully overlapping physical radio broadcast domain. But even in those cases where both are possible, WiND is preferable vs. ND because it reduces the need of broadcast (this is discussed in the introduction of [I-D.ietf-6lo-backbone-router]).

There are also numerous practical use cases in the wireless world where Links and Subnets are not P2P and not congruent:

- o IEEE std 802.11 infrastructure BSS enables one subnet per AP, and emulates a broadcast domain at L2. Infra ESS extends that and

recommends to use an IPv6 ND proxy [IEEEstd80211] to coexist with Ethernet connected nodes. WiND incorporates an ND proxy to serve that need and that was missing so far.

- o Bluetooth is Hub-and-Spoke at the MAC layer. It would make little sense to configure a different subnet between the central and each individual peripheral node (e.g., sensor). Rather, [RFC7668] allocates a prefix to the central node acting as router (6LR), and each peripheral host (acting as a host (6LR) forms one or more address(es) from that same prefix and registers it.
- o A typical Smartgrid networks puts together Route-Over MLSNs that comprise thousands of IPv6 nodes. The 6TiSCH architecture [I-D.ietf-6tisch-architecture] presents the Route-Over model over a [IEEEstd802154] Time-Slotted Channel-Hopping mesh, and generalizes it for multiple other applications. Each node in a Smartgrid network may have tens to a hundred others nodes in range. A key problem for the routing protocol is which other node(s) should this node peer with, because most of the possible peers do not provide added routing value. When both energy and bandwidth are constrained, talking to them is a bad idea and most of the possible P2P links are not even used. Peerings that are actually used come and go with the dynamics of radio signal propagation. It results that allocating prefixes to all the possible P2P Links and maintain as many addresses in all nodes is not even considered.

5.1. Case of LPWANs

LPWANs are by nature so constrained that the addresses and Subnets are fully pre-configured and operate as P2P or Hub-and-Spoke. This saves the steps of neighbor Discovery and enables a very efficient stateful compression of the IPv6 header.

5.2. Case of Infrastructure BSS and ESS

In contrast to IPv4, IPv6 enables a node to form multiple addresses, some of them temporary to elusive, and with a particular attention paid to privacy. Addresses may be formed and deprecated asynchronously to the association. Even if the knowledge of IPv6 addresses used by a STA can be obtained by snooping protocols such as IPv6 ND and DHCPv6, or by observing data traffic sourced at the STA, such methods provide only an imperfect knowledge of the state of the STA at the AP. This may result in a loss of connectivity for some IPv6 addresses, in particular for addresses rarely used and in a situation of mobility. This may also result in undesirable remanent state in the AP when a STA ceases to use an IPv6 address. It results

that snooping protocols is not a recommended technique and that it should only be used as last resort.

The recommended alternate is to use the IPv6 Registration method specified in p. By that method, the AP exposes its capability to proxy ND to the STA in Router Advertisement messages. In turn, the STA may request proxy ND services from the AP for one or more IPv6 addresses, using an Address Registration Option. The Registration state has a lifetime that limits unwanted state remanence in the network. The registration is optionally secured using [I-D.ietf-6lo-ap-nd] to prevent address theft and impersonation. The registration carries a sequence number, which enables a fast mobility without a loss of connectivity.

The ESS mode requires a proxy ND operation at the AP. The proxy ND operation must cover Duplicate Address Detection, Neighbor Unreachability Detection, Address Resolution and Address Mobility to transfer a role of ND proxy to the AP where a STA is associated following the mobility of the STA. The proxy ND specification associated to the address registration is [I-D.ietf-6lo-backbone-router]. With that specification, the AP participates to the protocol as a Backbone Router, typically operating as a bridging proxy though the routing proxy operation is also possible. As a bridging proxy, the proxy replies to NS lookups with the MAC address of the STA, and then bridges packets to the STA normally; as a routing proxy, it replies with its own MAC address and then routes to the STA at the IP layer. The routing proxy reduces the need to expose the MAC address of the STA on the wired side, for a better stability and scalability of the bridged fabric.

5.3. Case of Mesh Under Technologies

The Mesh-Under provides a broadcast domain emulation with reflexive and Transitive properties and defines a transit Link for IPv6 operations. It results that the model for IPv6 operation is similar to that of a BSS, with the root of the mesh operating an Access Point does in a BSS/ESS. While it is still possible to operate IPv6 ND, the inefficiencies of the flooding operation make the IPv6 ND operations even less desirable than in a BSS, and the use of WiND is highly recommended.

5.4. Case of DMC radios

IPv6 over DMC radios uses P2P Links that can be formed and maintained when a pair of DMC radios transmitters are in range from one another.

5.4.1. Using IPv6 ND only

DMC radios do not provide MAC level broadcast emulation. An example of that is OCB (outside the context of a BSS), which uses IEEE Std. 802.11 transmissions but does not provide the BSS functions.

It is possible to form P2P IP Links between each individual pairs of nodes and operate IPv6 ND over those Links with Link Local addresses. DAD must be performed for all addresses on all P2P IP Links.

If special deployment care is taken so that the physical broadcast domains of a collection of the nodes fully overlap, then it is also possible to build an IP Subnet within that collection of nodes and operate IPv6 ND.

The model can be stretched beyond the scope of IPv6 ND if an external mechanism avoids duplicate addresses and if the deployment ensures the connectivity between peers. This can be achieved for instance in a Hub-and-Spoke deployment if the Hub is the only router in the Subnet and the Prefix is advertised as not onlink.

5.4.2. Using Wireless ND

Though this can be achieved with IPv6 ND, WiND is the recommended approach since it uses more unicast communications which are more reliable and less impacting for other users of the medium.

Router and Hosts respectively send a compressed RA/NA with a SLLAO at a regular period. The period can be indicated in a RA as in an RA-Interval Option [RFC6275]. If available, the message can be transported in a compressed form in a beacon, e.g., in OCB Basic Safety Messages (BSM) that are nominally sent every 100ms. An active beaconing mode is possible whereby the Host sends broadcast RS messages to which a router can answer with a unicast RA.

A router that has Internet connectivity and is willing to serve as an Internet Access may advertise itself as a default router [RFC4191] in its RA. The NA/RA is sent over an Unspecified Link where it does not conflict to anyone, so DAD is not necessary at that stage.

The receiver instantiates a Link where the sender's address is not a duplicate. To achieve this, it forms an LLA that does not conflict with that of the sender and registers to the sender using [RFC8505]. If the sender sent an RA(PIO) the receiver can also autoconfigure an address from the advertised prefix and register it.




Figure 1: Initial Registration Flow

The lifetime in the registration should start with a small value ($X=R_{Min}$, TBD), and exponentially grow with each reregistration to a larger value ($X=R_{max}$, TBD). The IP Link is considered down when ($X=NbBeacons$, TBD) expected messages are not received in a row. It must be noted that the Link flapping does not affect the state of the registration and when a Link comes back up, the active -lifetime not elapsed- registrations are still usable. Packets should be held or destroyed when the Link is down.

P2P Links may be federated in Hub-and-Spoke and then in Route-Over MLSNs as described above. More details on the operation of WiND and RPL over the MLSN can be found in section 3.1, 3.2, 4.1 and 4.2.2 of [I-D.ietf-6tisch-architecture].

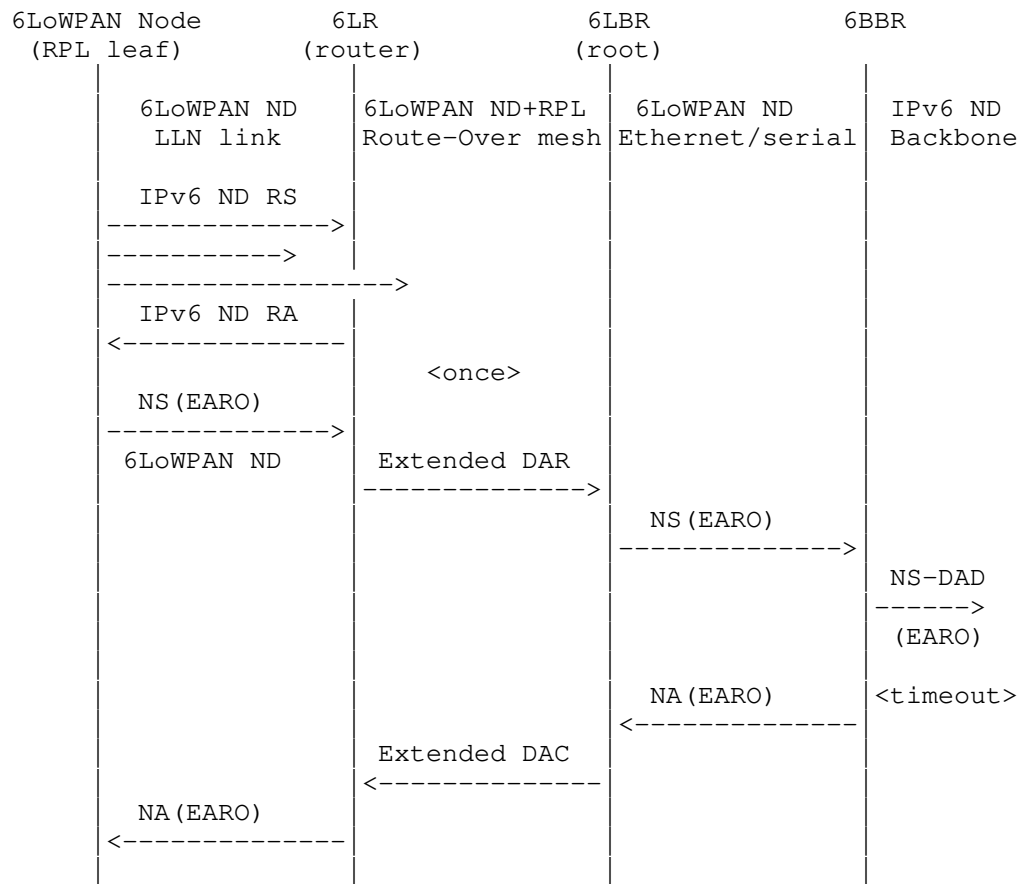


Figure 2: Initial Registration Flow over Multi-Link Subnet

An example Hub-and-Spoke is an OCB Road-Side Unit (RSU) that owns a prefix, provides Internet connectivity using that prefix to On-Board Units (OBUs) within its physical broadcast domain. An example of Route-Over MLSN is a collection of cars in a parking lot operating RPL to extend the connectivity provided by the RSU beyond its physical broadcast domain. Cars may then operate NEMO [RFC3963] for their own prefix using their address derived from the prefix of the RSU as CareOf Address.

6. IANA Considerations

This specification does not require IANA action.

7. Security Considerations

This specification refers to the security sections of IPv6 ND and WiND, respectively.

8. Acknowledgments

Many thanks to the participants of the 6lo WG where a lot of the work discussed here happened. Also ROLL, 6TiSCH, and 6LoWPAN.

9. References

9.1. Normative 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-12 (work in
progress), April 2019.
- [I-D.ietf-6lo-backbone-router]
Thubert, P., Perkins, C., and E. Levy-Abegnoli, "IPv6
Backbone Router", draft-ietf-6lo-backbone-router-11 (work
in progress), February 2019.
- [RFC3963] Devarapalli, V., Wakikawa, R., Petrescu, A., and P.
Thubert, "Network Mobility (NEMO) Basic Support Protocol",
RFC 3963, DOI 10.17487/RFC3963, January 2005,
<<https://www.rfc-editor.org/info/rfc3963>>.
- [RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and
More-Specific Routes", RFC 4191, DOI 10.17487/RFC4191,
November 2005, <<https://www.rfc-editor.org/info/rfc4191>>.
- [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>>.
- [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>>.

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

9.2. Informative References

- [I-D.bi-savi-wlan]
Bi, J., Wu, J., Wang, Y., and T. Lin, "A SAVI Solution for WLAN", draft-bi-savi-wlan-16 (work in progress), November 2018.
- [I-D.ietf-6tisch-architecture]
Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", draft-ietf-6tisch-architecture-20 (work in progress), March 2019.
- [I-D.ietf-mboned-ieee802-mcast-problems]
Perkins, C., McBride, M., Stanley, D., Kumari, W., and J. Zuniga, "Multicast Considerations over IEEE 802 Wireless Media", draft-ietf-mboned-ieee802-mcast-problems-05 (work in progress), April 2019.
- [I-D.ietf-rift-rift]
Team, T., "RIFT: Routing in Fat Trees", draft-ietf-rift-rift-05 (work in progress), April 2019.
- [I-D.thubert-6lo-unicast-lookup]
Thubert, P. and E. Levy-Abegnoli, "IPv6 Neighbor Discovery Unicast Lookup", draft-thubert-6lo-unicast-lookup-00 (work in progress), January 2019.
- [I-D.thubert-roll-unaware-leaves]
Thubert, P., "Routing for RPL Leaves", draft-thubert-roll-unaware-leaves-07 (work in progress), April 2019.
- [I-D.vyncke-6man-mcast-not-efficient]
Vyncke, E., Thubert, P., Levy-Abegnoli, E., and A. Yourtchenko, "Why Network-Layer Multicast is Not Always Efficient At Datalink Layer", draft-vyncke-6man-mcast-not-efficient-01 (work in progress), February 2014.

- [I-D.yourtchenko-6man-dad-issues]
Yourtchenko, A. and E. Nordmark, "A survey of issues related to IPv6 Duplicate Address Detection", draft-yourtchenko-6man-dad-issues-01 (work in progress), March 2015.
- [IEEE802154]
IEEE standard for Information Technology, "IEEE Std. 802.15.4, Part. 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks".
- [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)".
- [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>>.
- [RFC4903] Thaler, D., "Multi-Link Subnet Issues", RFC 4903, DOI 10.17487/RFC4903, June 2007, <<https://www.rfc-editor.org/info/rfc4903>>.

- [RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <<https://www.rfc-editor.org/info/rfc6550>>.
- [RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <<https://www.rfc-editor.org/info/rfc6775>>.
- [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>>.
- [RFC8273] Brzozowski, J. and G. Van de Velde, "Unique IPv6 Prefix per Host", RFC 8273, DOI 10.17487/RFC8273, December 2017, <<https://www.rfc-editor.org/info/rfc8273>>.

Author's Address

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

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