

INTERNET-DRAFT
Intended Status: Informational
Expires: December 31, 2018

Hudson Ayers
Paul Crews
Hubert Teo
Conor McAvity
Amit Levy
Philip Levis
Stanford University
June 29, 2018

Design Considerations For Low Power Internet Protocols
draft-ayers-low-power-interop-00

Abstract

This document discusses guidelines for specifying low-power Internet protocols in order to improve implementation interoperability. These guidelines are based around the importance of balancing memory usage and energy efficiency, and the importance of not relying on Postel's law when dealing with low resource devices. This document applies these guidelines to the IPv6 over low-power wireless personal area networks (6LoWPAN) Internet Standard, suggesting changes that would make it more likely for implementations to interoperate.

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

This Internet-Draft will expire on <Expiry Date>

Copyright and License Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1	Introduction	3
1.1	Terminology	4
2.	6LoWPAN Interoperability Study	4
2.1	Incomplete Implementations	5
2.2	Unrealistic Bounds	6
2.2.1	Maximum Header Decompression	6
2.2.2	Arbitrary Next Header Compression	6
2.3	No Pairing Interoperates	7
3	Implementation Concerns	9
3.1	Processor Resources	9
5	Contributing Factors	11
6	Design Guidelines	11
6.1	Guideline 1: Capability Spectrum	11
6.1.1	Guideline 1 Application to 6LoWPAN	12
6.2	Guideline 2: Capability Discovery	13
6.2.1	Guideline 2 Application to 6LoWPAN	13
6.3	Guideline 3: Provide Reasonable Bounds	14
6.3.1	Guideline 3 Application to 6LoWPAN	14
6.4	Guideline 4: Don't Break Layering	15
6.4.1	Guideline 4 Application to 6LoWPAN	16
7	Security Considerations	18
8	IANA Considerations	18
9	References	18
9.1	Normative References	18
9.2	Informative References	19
	Authors' Addresses	20

1 Introduction

Interoperability is critical for the Internet. Not only do edge-devices need to interoperate with the broader Internet, they should also interoperate with other devices in the same network. Historically, though, embedded systems and sensor networks have been vertical silos of proprietary technologies, each using custom network protocols and homogeneous implementations. Networks typically require specialized gateways and cannot easily include devices from different vendors for a variety of applications.

Interoperability is just as important in the Internet of Things, but end hosts in the Internet of Things are less resourceful, more diverse in capability, and less well audited than typical Internet end hosts. Using IP allows devices from different manufacturers, running completely different software stacks, to interoperate, share services, and compose into larger, more complex applications. This interoperability should exist not only between IoT devices communicating with hosts across the broader Internet, but also between IoT devices in the same low power wireless network. The presence of such interoperability precludes the need for multiple gateways to support different devices, simplifies network management, and allows for efficient, logical communication between nearby devices.

To address this problem, the IETF published a series of RFCs detailing a standard format for transmitting IPv6 packets over low-power wireless link layers such as IEEE 802.15.4 {RFC 4919}[RFC 4944][RFC 6282][RFC 6775]. The 6LoWPAN RFCs define a fragmentation format, a compression format, and more. These 6LoWPAN standards have been adopted by a number of popular embedded operating systems, including Contiki {CONTIKI}, RIOTOS {RIOT}, OpenThread [OPENTHREAD], mbedOS [ARM], and TinyOS [TINYOS].

Unfortunately, none of these implementations are complete. Each implementation supports different subsets of 6LoWPAN. As a result, devices built using different embedded operating systems cannot interoperate. In fact, for every possible pairing, one implementation is likely to transmit 6LoWPAN packets which the other cannot process.

This paper explores the reasons behind the lack of interoperability in practice, and argues that this results from the protocol too heavily prioritizing radio efficiency over processor resources, and failing to consider the broad range of devices which embedded operating systems will attempt to support. This document proposes four guidelines for designing interoperable protocols for low-power wireless networks, and explain them through an example application to two 6LoWPAN standards - RFC 4944 and RFC 6282. These guidelines are

informed by an empirical analysis of existing 6LoWPAN implementations as well as experience implementing a full 6LoWPAN stack for the Tock operating system [TOCK].

1.1 Terminology

Readers are expected to be familiar with all terms and concepts discussed in "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals", "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks".

Readers would benefit from reading 6LoWPAN Neighbor Discovery (ND), 6LoWPAN routing requirements, and 6LoWPAN design/application spaces for additional details of 6LoWPAN work.

2. 6LoWPAN Interoperability Study

The IETF's 6LoWPAN working group has been concerned with interoperability between implementation since inception [6LO-CHART]. Indeed, members of the working group have organized ``Plugtests'', where vendors verified correct implementation of the 6LoWPAN specifications and tested interoperability with other vendors. Unfortunately, detailed results of these plugtests are not publicly available.

A node sending IPv6 packets using 6LoWPAN may fragment the packet or compress headers in a large number of ways permitted by the specification. These choices depend both on the properties of the packet (e.g. whether it is a UDP packet or if the origin and destinations are in the same subnet) as well as on which compression and fragmentation options the sender chooses to use. For two nodes to interoperate, the 6LoWPAN implementation on each node must be able to receive any packet the other node might send.

This document investigates the interoperability of 6LoWPAN implementations from five common embedded software platforms: Riot OS, Contiki, OpenThread, TinyOS, and ARM Mbed. This study is concerned, specifically, with each implementation's ability to receive and decode 6LoWPAN packets sent from other implementations. This document does not explore whether devices built using these implementations could form a network in the first place, since 6LoWPAN leaves much of the network formation process (e.g. discovery and joining) unspecified.

The compatibility analysis in this document, between five common

embedded software frameworks with 6LoWPAN implementations, is based in three discoveries. First, a determination of how completely each implementation implements the 6LoWPAN specification was obtained by directly examining their source code. Second, this code analysis was extended to verify that under some circumstances, each implementation sends packets using compression or fragmentation options which another implementation cannot decode. Finally, it was discovered that even in cases where two implementations use compatible compression and fragmentation options, different implementation choices, such as header decompression bounds, limit their interoperability.

As a result, no pairing of these five implementations is fully interoperable.

2.1 Incomplete Implementations

The 6LoWPAN protocol consists of a large number of complex and mostly independent features which use the link-layer frame efficiently via compression and fragmentation optimizations. Examining the code and documentation for each of the aforementioned 6LoWPAN stacks reveals that the stacks do not uniformly implement the specification. In fact, each specification implements a different subset of the requirements in the 6LoWPAN specification, and none implements the entire specification to the letter. A visualization of the mismatched feature support across different 6LoWPAN implementations can be found in Table 1. Note that OT is an abbreviation used throughout this document to refer to OpenThread.

Feature	Contiki	OT	Riot	Arm	TinyOS
Uncompressed IPv6	o		o	o	o
6LoWPAN Fragmentation	o	o	o	o	o
1280 byte packets	o	o	o	o	o
Dispatch_IPHC header prefix	o	o	o	o	o
IPv6 Stateless Address Compression	o	o	o	o	o
Stateless multicast address compression	o	o	o	o	o
802.15.4 16 bit short address support		o	o	o	o
IPv6 Address Autoconfiguration	o	o	o	o	o
IPv6 Stateful Address Compression	o	o	o	o	o
Stateful multicast address compression		o	o	o	
IPv6 TC and Flow label compression	o	o	o	o	o
IPv6 NH Compression: Tunneled IPv6		o		o	o
IPv6 NH Compression: UDP	o	o	o	o	o
UDP port compression	o	o	o	o	o
UDP checksum elision					o

Compression + headers past first frag						
-----	-----	-----	-----	-----	-----	-----

Table 1: 6LoWPAN Interoperability Matrix

2.2 Unrealistic Bounds

Beyond the variation in what portions of the 6LoWPAN specification each stack implements, there also exists significant variation in how each stack handles certain implementation-specific details. Some of these details have little impact on interoperability, such as decisions regarding how many fragments a stack holds for a given packet before dropping all of them, whether to allow for reconstruction of multiple packets simultaneously, and how long to hold onto fragments for which the rest of the packet has not yet arrived. Other details, however, differ in ways that significantly affect interoperability between stacks. A discussion of two such details follows:

2.2.1 Maximum Header Decompression

Each of the stacks analyzed imposes some limit on the maximum amount of header decompression possible for a received packet. Such a limit is necessary to ensure that packet and fragment buffers within a stack are large enough for received packets. The maximum amount of header decompression allowed by the 6LoWPAN specification is about 1200 bytes, basically, if an entire MSS IPv6 packet was sent containing only compressed headers. Some of the stacks analyzed decompress fragments directly into the MSS buffer which will eventually contain the entire IPv6 packet, and thus support this bound. Other stacks impose significantly lower limits - limits low enough that packets could easily be constructed within the 6LoWPAN specification that would exceed these limits. For example, Contiki's limit of 38 bytes of header decompression is exceeded by any packet for which the IP header is maximally compressed (38 bytes) and the UDP header is compressed at all. Accordingly, certain stacks would send packets with a significant amount of header compression, but other stacks would silently drop these packets due to lacking buffer space for fragments requiring that much decompression. Furthermore, these stacks do not given any indication back to the sender that a packet has been dropped for this reason, making it difficult for the sending stack to identify how to adjust its transmission to successfully deliver data.

2.2.2 Arbitrary Next Header Compression

Several of the 6LoWPAN stacks also impose limits on the arbitrary

compression/decompression of IPv6 extension headers and next headers required by the specification. The headers which must be handled are as follows:

- IPv6 Hop-By-Hop Options Header
- IPv6 Routing Header
- IPv6 Fragment Header
- IPv6 Destination Options Header
- IPv6 Mobility Header
- IPv6 Next Header
- UDP Next Header

Further, 6LoWPAN implementations are expected to be able to decompress at least one of each of these headers, and up to two Destination Options headers, in almost any order. Handling all of these possible cases can result in complex state machines, convoluted code, and increase in code size and RAM use. Therefore, several of the stacks examined impose a limit on this arbitrary next header decompression - namely, Contiki and Riot. Both of these stacks only check for the UDP Next Header. This greatly simplifies the code required for decompression of next headers in these stacks as compared to the others, which require recursion to handle this arbitrary compression. The offshoot of this simplified code, however, is that these stacks will drop packets with certain compressed extension header configurations when other stacks send such messages.

2.3 No Pairing Interoperates

These interoperability concerns are more than theoretical: existing 6LoWPAN stacks generate valid packets that other stacks discard. This proves that missing receive functionality is not simply a case of limited 6LoWPAN stacks abstaining from handling packets which no existing stacks ever generate.

What follows is a listing of each of the 10 possible combinations of 6LoWPAN stacks, accompanied by a single example packet which can be generated by one of the stacks in the pairing which the other stack would not receive.

Contiki, OpenThread : Contiki generated message using uncompressed IPv6

Contiki, Riot: Riot generated message using stateful multicast address compression

Contiki, Mbed: Mbed generated message using compressed, tunneled IPv6

Contiki, TinyOS: TinyOS generated message containing compressed IPv6 extension headers

OpenThread, Riot: OpenThread generated message containing any of the IPv6 extension headers, which the OpenThread stack automatically compresses

OpenThread, Mbed: Mbed generated IPv6 packet containing the IPv6 mobility header

OpenThread, TinyOS: OpenThread generated message for which the destination address is compressed using stateful multicast compression

Riot, Mbed: Mbed generated IPv6 message containing any compressed next header other than the UDP header

Riot, TinyOS: Riot generated message for which the destination address is compressed using stateful multicast compression

Mbed, TinyOS: Mbed generated Neighbor Discovery message using the 6LoWPAN context option as specified in RFC 6775.

This is a non-exhaustive listing, and for most of these pairings several message formats exist which could be generated by one that would be dropped by the other. Each instance for which a claim is made that packets could be easily generated has been verified via code analysis.

In addition to this code analysis, tests were performed to present further evidence that several of these packets formats could easily be generated via typical use of these 6LoWPAN stacks.

These tests involved slightly modifying basic example networking apps on each stack, such that the existing 6LoWPAN interface could be used to send certain packets. These modified examples were flashed onto embedded hardware platforms supported by each. The transmitted packets were captured using a wireless packet sniffer, and the sniffed packets analyzed using Wireshark. This exercise verified that these non-interoperable packets could in fact be sent. Further description of this hardware generation of select packets can be found in [DESIGN].

3 Implementation Concerns

The 6LoWPAN specification was created with a clear goal---to allow for IPv6 connectivity over a link-layer with an order of magnitude smaller frame sizes than Ethernet. Unfortunately, fragmented IPv6 on its own requires header overhead much greater than typical wireless protocols designed for low power devices. As a result, the specification places an extreme focus on minimizing protocol overhead and, thus, radio utilization.

The primary problem with 6LoWPAN is that this focus was taken too far. This focus has resulted in complex implementations that require significant processor resources. In order for devices to interoperate, they must be able to parse any valid received 6LoWPAN packet that might be sent by others.

In practice, many 6LoWPAN implementations do not implement the entire specification and, therefore, are not interoperable. This is not a result of poor software design, but rather intentional choices to implement different subsets of the specification that favor limited RAM and code size, security concerns, and minimizing engineering effort.

In fact, in some cases even these incomplete 6LoWPAN implementations systems are too resource intensive for some devices. As a result, several implementations allow the developer to remove portions of the 6LoWPAN stack during compilation. Even when implementations use overlapping portions of the specification, additional interoperability conflicts arise from different choices of memory bounds for decompression.

3.1 Processor Resources

Evidence that developers of these 6LoWPAN stacks were concerned about 6LoWPAN's consumption of processor resources is baked into the design of each. One of the primary indicators that each implementation was concerned with code size is the prevalence of options to compile limited subsets of the 6LoWPAN stack. For example, Contiki defines the `SICSLOWPAN_CONF_COMPRESSION` compilation flag, which can be set to force all Contiki packets (sent and received!) to be processed as uncompressed IPv6. Riot presents extensive compilation options for 6LoWPAN, allowing for the exclusion of all IPHC compression, the exclusion of context based compression alone, the exclusion of fragmentation, the exclusion of ND, and the exclusion of next header compression. The Mbed stack allows users to exclude elements of the IPv6 stack such as security features, routing specific features, link-layer features, and more. Further, Mbed defines macros which can be used to save RAM at the expense of flash, or vice-versa. TinyOS by

default removes all code in a stack that is not being used by an application, and this can easily be observed by compiling different 6LoWPAN application binaries.

Table 2 shows the code size overhead of each of the five implementations broken into independent overheads for compression, fragmentation, mesh and broadcast headers, as well as totals for 6LoWPAN and the entire networking stack including physical layer drivers, IPv6, UDP, ICMP, etc.

Stack	Code Size Measurements (Bytes)				
	IP-All	6Lo-All	Compression	Frag	Mesh/Bcast Hdr
Contiki	37538	11262	5952	3319	N/A
OT	42262	26375	4146-20000	1310	4500
Riot	30942	7500	>4712	1514	N/A
Arm Mbed	46030	22061	17900	3104	1331
TinyOS	37312	16174	----	----	600

Table 2: 6LoWPAN Stack Code Size

The methodology use to collect these values can be found in [DESIGN]. These results likely overestimate the overhead of fragmentation and underestimate the overhead of certain kinds of compression since some of the complexity of compression is born on the fragmentation logic. Moreover, for OpenThread and Arm Mbed, which required manual examination of binaries, the results almost certainly underestimate the overhead of all 6LoWPAN components since we only counted procedures which unambiguously implemented particular functionality, though some of the complexity is implemented in other portions of the stack. In summary:

- 6LoWPAN stack developers were concerned with processor resource requirements of the protocol.
- Fragmentation, the only portion of 6LoWPAN that's strictly necessary for sending IPv6 packets, consumes significantly less ROM than compression.
- Implementations with more complete adherence to compression specification consume more code for compression
- Mesh and broadcast headers are relatively expensive given that few real-world applications use them

5 Contributing Factors

Several fundamental factors contributed to 6LoWPAN's interoperability problems.

When writing low power networking specifications, an important "slider" exists - the tradeoff between code size and protocol efficiency. This tradeoff is similar to the historically significant tradeoff between RAM and code size. Techniques such as advanced MAC and physical layers, and tracking the state of a network can reduce packet sizes and, thus, radio energy consumption. However, these techniques typically require larger and more complex implementations.

Even moving beyond the constraints of code size, added complexity harms interoperability in the general case, and complex implementations are undesirable in the space of low power embedded devices. Finally, 6LoWPAN failed to consider the reality that some implementations of the protocol may be incomplete, and accordingly failed to include any affirmative indications of interoperability failures, with interoperability failures instead only being visible as silent packet drops. All of these factors contributed to 6LoWPAN's interoperability problems, and inspire the guidelines that follow.

6 Design Guidelines

This section describes four protocol design guidelines which, if followed, lead to low-power protocols that are more likely to have interoperable implementations. In the next section, these are further explained by showing how each can be applied to 6LoWPAN.

6.1 Guideline 1: Capability Spectrum

A low power protocol should be implementable on devices which are at the low end of code and RAM resources. Rather than require every device pay the potential energy costs of fewer optimizations, a protocol should support a spectrum of device capabilities. This spectrum defines a clear ordering via which especially resource constrained devices can reduce code size or RAM use by eliding features. Such a spectrum makes a protocol usable by extremely low resource devices without forcing more resourceful devices to communicate inefficiently.

This capability spectrum should be a linear scale. For a device to support capability level N, it must also support all lower capability levels. More complex configuration approaches (e.g., a set of independent options) would allow for a particular application or implementation to be more efficient, picking the features that give

the most benefit at the least complexity cost. However, this sort of optimization then makes interoperability more difficult, as two devices must negotiate which features to use.

6.1.1 Guideline 1 Application to 6LoWPAN

Application of this guideline would require replacing the large collection of "MUST" requirements - those "features" in Table 1 - into 6 levels of functionality. These levels prioritize features that provide the best packet size savings given the resulting implementation complexity. For example, the greatest savings results from compressing 128-bit IPv6 addresses.

0. Uncompressed IPv6

0a. Uncompressed IPv6

0b. 6LoWPAN Fragmentation and the Fragment Header

0c. 1280 Byte Packets

1. IPv6 Compression Basics + Stateless Address Compression

1a. Support for the Dispatch_IPHC Header Prefix

1b. Correctly handle elision of IPv6 length and version

1c. Stateless compression of unicast addresses

1d. Stateless compression of multicast addresses

1e. Compression even when 16 bit addresses are used at the link layer

1f. IPv6 address autoconfiguration

2. Stateful IPv6 Address Compression

2a. Stateful compression of unicast addresses

2b. Stateful compression of multicast addresses

3. IPv6 Traffic Class and Flow Label Compression

3a. Traffic Class compression	3b. Flow Label Compression	3c. Hop Limit Compression
-------------------------------	----------------------------	---------------------------

4. IPv6 and UDP NH Compression + UDP Port Compression

- 4a. Handle Tunneled IPv6 correctly
 - 4b. Handle the compression of the UDP Next Header
 - 4c. Correctly handle elision of the UDP length field
 - 4d. Correctly handle the compression of UDP ports
 - 4e. Correctly handle messages for which headers go on longer than the first fragment, and the headers in the first fragment are compressed.
5. Entire Specification
- 5a. Support the broadcast header and the mesh header as described in RFC 4944
 - 5b Support compression of all IPv6 Extension headers

The classes in this scale do not precisely reflect the current feature support of the implementations described above. For example, Contiki supports UDP port compression (level 5) but does not support 802.15.4 short addresses (level 2) or tunneled IPv6 (level 5): following this formulation, Contiki only provides level 1 support. If Contiki supported 16-bit addresses, it would provide level 4 support.

The specific spectrum presented here is based off of measurements of code size, the saved bits that each additional level of compression allows for, and observations of existing 6LoWPAN implementations.

6.2 Guideline 2: Capability Discovery

The second guideline immediately follows from the first: if two implementations may have different capability levels, there should be an explicit mechanism by which two devices can efficiently discover what level to use when they communicate

If two devices wish to communicate, they default to the lower of their supported capability levels. For example, suppose a TinyOS device supports level 2 and a Contiki device supports level 4; Contiki must operate at level 2 when communicating with the TinyOS device. This requires keeping only a few bits of state for any device to communicate with. Also, note that this state is per-hop; for a layer 3 protocol like IP, it is stored for link-layer neighbors (not IP endpoints) and so does not require knowledge of the whole network.

6.2.1 Guideline 2 Application to 6LoWPAN

6lowpan could implement capability discovery using two mechanisms: neighbor discovery (ND) and ICMP. Neighbor discovery allows devices to probe and determine capability levels, while ICMP allows devices to determine when incompatible features are used, or when ND is not available.

Neighbor discovery: 6LoWPAN ND should add an option that allows a device to communicate its capability class during association with a network. The inclusion of a few extra bits in ND messages would allow all devices that learn neighbor addresses via ND to also know how to send packets which that neighbor can receive. This option minimizes the energy cost of communicating capabilities. It is worth noting that [RFC7400] already employs a similar method for communicating whether devices implement General Header Compression: adding such an option is clearly viable.

ICMP: All IPv6 devices are already required to support ICMP. A new ICMPv6 message type - 6LoWPAN Class Unsupported - should be added, which could be sent in response to messages received encoded using a 6LoWPAN class higher than the class of the receiving host. This would allow for communication of capabilities even in networks not constructed using IPv6 ND. This ICMPv6 message would allow hosts to indicate exactly what class the receiving host does support, preventing any need for repeated retransmissions using different compression or fragmentation formats.

6.3 Guideline 3: Provide Reasonable Bounds

Specifications should impose reasonable bounds on recursive or variable features so implementations can bound RAM use. These bounds have two benefits. First, it allows implementations to safely limit their RAM use without silent interoperability failures. E.g., today, if an mbed device sends a 6lowpan packet whose compression is greater than 38 bytes to a Contiki device, Contiki will silently drop the packet. Second, it ensures that capability discovery is sufficient to interoperate.

The original designers of a specification may not know exactly what these values should be. This is not a new problem: TCP congestion control, for example, had to specify initial congestion window values. The bounds should initially be very conservative. Over time, if increasing resources or knowledge suggests they should grow, then future devices will have the onus of using fewer resources to interoperate with earlier ones.

6.3.1 Guideline 3 Application to 6LoWPAN

Section 2 discussed two unreasonable bounds which affect 6LoWPAN

interoperability. The first is the 1280 byte bound on maximum header decompression (the amount a header will grow when decompressed). A bound allows implementations to conserve RAM. As a result, some implementations impose their own lower bounds, but these bounds do not agree so some stacks cannot decompress some packets sent by other stacks. The lack of a bound on arbitrary next header compression was demonstrated as adding significant complexity to implementations to service packets which should rarely be used.

To address this, maximum header decompression in 6LoWPAN packets should be bounded to 50 bytes. This bound allows for significant RAM savings in implementations that decompress first fragments into the same buffer in which the fragment was originally held prior to any copying into a 1280 byte buffer.

Second, the requirement for compression of interior headers for tunneled IPv6 should be removed. Currently, section 4.2 of RFC 6282 states "When the identified next header is an IPv6 Header...The following bytes MUST be encoded using LOWPAN_IPHC". This is problematic because it places no bound on how many tunneled IPv6 headers may need to be compressed or decompressed, creating locations in code that require unbounded amounts of recursion. Implementations should adjust their path MTU constraints and responses to support inserting source routing headers, rather than tunnel IPv6.

This change would limit the complexity of arbitrary next header compression slightly. In addition, an ordering should be imposed on the order of IPv6 extension options if they are to be compressed. This would allow for implementations to avoid recursive functions to decompress these headers, and instead use simple if/else statements. If for some reason IPv6 extension headers must be placed in a different order for a particular packet, those options must be sent uncompressed.

6.4 Guideline 4: Don't Break Layering

Designers should ensure that interoperability is a central priority for specifications throughout the design process, and that interoperability is not simply assumed from the fact that devices will be communicating via a shared protocol. In particular, specifications should be careful that considerations introduced to save energy in certain scenarios should not make assumptions about the rest of the stack. Layering is a foundational network design principle. As the difficulty NATs introduced to Internet connectivity in the early 2000s demonstrated, breaking layering can introduce unforeseen and extremely difficult to fix interoperability problems.

The appeal of cross-layer optimization in embedded systems is even

stronger than in traditional computers. Designed for a specific application, a developer can understand and know exactly how the entire system works, from hardware to application code. However, while this whole-system knowledge makes sense for a particular device or iteration of an application, long-lived systems will evolve and change. This is especially true if the device will need to interoperate with new gateways or application devices. Furthermore, as embedded systems have grown more complex, their software has begun to resemble more traditional systems. Rather than write software from scratch every time, systems use and draw on existing operating systems as well as libraries. By breaking layering, cross-layer optimizations require that developers own and customize the entire software stack.

6.4.1 Guideline 4 Application to 6LoWPAN

UDP checksum compression, as defined in section 4.3.2 of RFC 6282, should be removed from the 6LoWPAN specification. The RFC says that a higher layer may request the checksum be elided if it has an integrity mechanism that covers the UDP header. At first glance, this seems sufficient: if the UDP header is covered by a message integrity code (MIC) or other checksum, then corrupted packets will be correctly dropped.

However, it misses an important error case: if the UDP ports are corrupted, then a packet missing a checksum may be delivered to the wrong application, and this incorrect application may not impose a replacement integrity measure or know one exists. It therefore cannot verify the MIC. Furthermore, protecting the header with a link-layer MIC is insufficient, as it only protects packets against sub-link corruption.

The end-to-end principle [E2E], foundational to all modern network design, says that only endpoints can verify correct communication. The only place that can safely verify the UDP header is the UDP stack. It is worth noting that the seminal example that led to definition of the end-to-end principle was a memory corruption: packets held in memory to be sent were corrupted before being sent. The recommended workarounds in RFC 6282 are vulnerable to such an event. A packet sent by an application that elides the UDP checksum could be corrupted in memory before the link-layer MIC is computed. Such a packet would be successfully received by the destination and dispatched to the wrong application, which would not check the application-level MIC.

The payoff of UDP checksum compression is not even significant - 2 bytes of checksum is a small portion of a 127 byte frame. The problematic nature of UDP checksum compression is further

demonstrated by the fact that only one of the five stacks analyzed in this document implements the feature.

7 Security Considerations

This informational document does have some implications for security if followed.

First, capability advertisements of the type recommended in this document are liable to leak some information regarding the type of device sending those advertisements. In any situation for which this information is privileged, such advertisements must be suppressed.

Second, implementations should be careful not to take for granted that the suggestions in this document will be implemented by all other transmitting devices. Accordingly, though this document recommends reasonable bounds, receivers still must be careful to prevent buffer overflows in the event these bounds are not followed.

Finally, it is worth noting that breaking layering has clear security implications, and that the recommendation in this document to avoid this practice should be expected to improve security by allowing the security protocols in place at individual layers to work as intended.

8 IANA Considerations

This is an informational document, and accordingly does not formally request any IANA changes. However, it is worth noting that the example application of the guidelines to 6LoWPAN would require some changes by IANA, if actually implemented.

Namely, IANA would be requested to update some of the "6LoWPAN Capability Bits" under the "Internet Control Message Protocol Version 6 (ICMPv6) Parameters" registry such that some of the unassigned bits could be repurposed for capability advertisements as described in this document.

Additionally, IANA would be requested to update the "IPv6 Neighbor Discovery Option Formats" registry to include a new ND option format for capability advertisements [RFC4861].

9 References

9.1 Normative References

[RFC4919] Kushalnagar, N., Montenegro, G., and C. Schumacher, "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals", RFC

4919, DOI 10.17487/RFC4919, August 2007, <<https://www.rfc-editor.org/info/rfc4919>>.

- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, September 2007, <<http://www.rfc-editor.org/info/rfc4944>>.
- [RFC6282] Hui, J. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, September 2011, <<http://www.rfc-editor.org/info/rfc6282>>.
- [RFC6775] Shelby, Z., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, November 2012, <<http://www.rfc-editor.org/info/rfc6775>>.

9.2 Informative References

- [DESIGN] Ayers, H. et al., "Design Considerations for Low Power Internet Protocols", Arxiv, June 2018.
- [TINYOS] TinyOS Alliance, "TinyOS", 2018, <<https://github.com/tinyos/tinyos-main>>.
- [ARM] ARM Mbed, "ARM Mbed OS", 2018, <<https://github.com/ARMmbed/mbed-os>>.
- [RIOT] FU Berlin, "Riot OS", 2018, <<https://github.com/RIOT-OS/RIOT>>.
- [CONTIKI] Dunkels, A., "Contiki OS", 2018, <<https://github.com/contiki-os/contiki>>.
- [OPENTHREAD] Nest, "OpenThread", 2018, <<https://github.com/openthread/openthread>>.
- [TOCK] Levy, A., Campbell, B., Pannuto, P., Dutta, P., Levis, P., "The Case for Writing a Kernel in Rust", APSys, 2017, <<https://doi.org/10.1145/3124680.3124717>>.
- [6LO-CHART] Lemon, T., "IPv6 over Low power WPAN WG Charter", IETF, 2005, <<https://datatracker.ietf.org/doc/charter-ietf-6lowpan/>>.
- [E2E] Saltzer, J. H., Reed, D. P., Clark, D. D., "End-to-end Arguments in System Design", ACM Trans. Comput. Syst., November 1984.

[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>>.

Authors' Addresses

Hudson Ayers
Stanford University

EMail: hayers@stanford.edu

Paul Crews
Stanford University

EMail: ptcrews@stanford.edu

Hubert Hua Kian Teo
Stanford University

EMail: hteo@stanford.edu

Conor McAvity
Stanford University

EMail: cmcavity@stanford.edu

Amit Levy
Stanford University

EMail: levya@cs.stanford.edu

Philip Levis
Stanford University

EMail: pal@stanford.edu

6Lo Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 24, 2019

J. Hou
B. Liu
Huawei Technologies
Y-G. Hong
ETRI
X. Tang
SGEPRI
C. Perkins
Futurewei
October 21, 2018

Transmission of IPv6 Packets over PLC Networks
draft-hou-6lo-plc-05

Abstract

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

Status of This Memo

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

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

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

This Internet-Draft will expire on April 24, 2019.

Copyright Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

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

1. Introduction

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

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

The following sections provide a brief overview of PLC, then describe transmission of IPv6 packets over PLC networks. The general approach is to adapt elements of the 6LoWPAN specifications [RFC4944], [RFC6282], and [RFC6775] to constrained PLC networks. Compared to [I-D.popa-6lo-6loplc-ipv6-over-ieee19012-networks], this document provides a structured and greatly expanded specification of an adaptation layer for IPv6 over PLC (6LoPLC) networks.

2. Requirements Notation and Terminology

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

This document often uses the following acronyms and terminologies:

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

AMI: Advanced Metering Infrastructure

BBPLC: Broadband Power Line Communication

CID: Context ID

Coordinator: A device capable of relaying messages.

DAD: Duplicate Address Detection

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

EV: Electric Vehicle

IID: IPv6 Interface Identifier

IPHC: IP Header Compression

LAN: Local Area Network

MSDU: MAC Service Data Unit

MTU: Maximum Transmission Unit

NBPLC: Narrowband Power Line Communication

OFDM: Orthogonal Frequency Division Multiplexing

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

PLC: Power Line Communication

PSDU: PHY Service Data Unit

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

RA: Router Advertisement

WAN: Wide Area Network

The terminology used in this draft is aligned with IEEE 1901.2

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

Table 1: Terminology Mapping between PLC standards

3. Overview of PLC

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

3.1. Protocol Stack

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




Figure 1: PLC Protocol Stack

3.2. Addressing Modes

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

3.3. Maximum Transmission Unit

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

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

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

3.4. Routing Protocol

Routing protocols suitable for use in PLC networks include:

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

4. IPv6 over PLC

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

4.1. Stateless Address Autoconfiguration

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

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

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

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

16_bit_PAN:00FF:FE00:16_bit_short_address

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

YYYY:YYFF:FE00:0XXX

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

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

4.2. IPv6 Link Local Address

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




Figure 2: IPv6 Link Local Address for a PLC interface

4.3. Unicast Address Mapping

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

4.3.1. Unicast Address Mapping for IEEE 1901.1

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




Figure 3: Unicast Address Mapping for IEEE 1901.1

Option fields:

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

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

NID: 24-bit Network IDentifier

Padding: 12 zero bits

TEI: 12-bit Terminal Equipment Identifier

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

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

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




Figure 4: Unicast Address Mapping for IEEE 1901.2

Option fields:

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

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

PAN ID: 16-bit PAN Identifier

Padding: 16 zero bits

Short Address: 16-bit short address

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

4.4. Neighbor Discovery

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

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

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

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

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

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

4.5. Header Compression

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

4.6. Fragmentation and Reassembly

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

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

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

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

4.7. Extension at 6lo Adaptation Layer

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

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

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




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

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

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

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

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

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

5. Internet Connectivity Scenarios and Topologies

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

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




Figure 7: PLC Star Network connected to the Internet

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

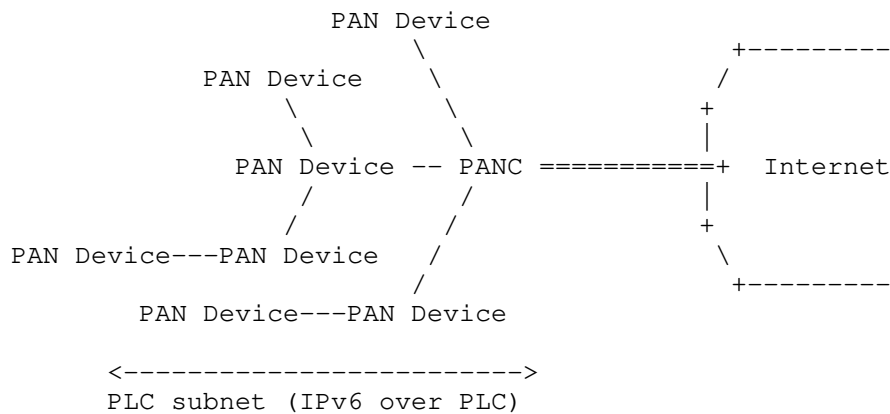


Figure 8: PLC Tree Network connected to the Internet

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




Figure 9: PLC Mesh Network connected to the Internet

6. IANA Considerations

There are no IANA considerations related to this document.

7. Security Consideration

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

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

8. Acknowledgements

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

9. References

9.1. Normative References

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

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

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

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

- [IEEE_1901.1]
IEEE-SA Standards Board, "Standard for Medium Frequency (less than 15 MHz) Power Line Communications for Smart Grid Applications", IEEE 1901.1, May 2018, <<http://sites.ieee.org/sagroups-1901-1>>.
- [IEEE_1901.2]
IEEE-SA Standards Board, "IEEE Standard for Low-Frequency (less than 500 kHz) Narrowband Power Line Communications for Smart Grid Applications", IEEE 1901.2, October 2013, <<https://standards.ieee.org/findstds/standard/1901.2-2013.html>>.
- [ITU-T_G.9903]
International Telecommunication Union, "Narrowband orthogonal frequency division multiplexing power line communication transceivers for G3-PLC networks", ITU-T G.9903, February 2014, <<https://www.itu.int/rec/T-REC-G.9903>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, DOI 10.17487/RFC2464, December 1998, <<https://www.rfc-editor.org/info/rfc2464>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<https://www.rfc-editor.org/info/rfc4944>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<https://www.rfc-editor.org/info/rfc6282>>.

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

9.2. Informative References

- [I-D.popa-6lo-6loplc-ipv6-over-ieee19012-networks] Popa, D. and J. Hui, "6LoPLC: Transmission of IPv6 Packets over IEEE 1901.2 Narrowband Powerline Communication Networks", draft-popa-6lo-6loplc-ipv6-over-ieee19012-networks-00 (work in progress), March 2014.
- [IEEE_1901.2a] IEEE-SA Standards Board, "IEEE Standard for Low-Frequency (less than 500 kHz) Narrowband Power Line Communications for Smart Grid Applications - Amendment 1", IEEE 1901.2a, September 2015, <<https://standards.ieee.org/findstds/standard/1901.2a-2015.html>>.
- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <<https://www.rfc-editor.org/info/rfc3315>>.
- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, DOI 10.17487/RFC3972, March 2005, <<https://www.rfc-editor.org/info/rfc3972>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.
- [RFC4429] Moore, N., "Optimistic Duplicate Address Detection (DAD) for IPv6", RFC 4429, DOI 10.17487/RFC4429, April 2006, <<https://www.rfc-editor.org/info/rfc4429>>.

- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, DOI 10.17487/RFC4941, September 2007, <<https://www.rfc-editor.org/info/rfc4941>>.
- [RFC5535] Bagnulo, M., "Hash-Based Addresses (HBA)", RFC 5535, DOI 10.17487/RFC5535, June 2009, <<https://www.rfc-editor.org/info/rfc5535>>.
- [RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <<https://www.rfc-editor.org/info/rfc7217>>.
- [RFC8036] Cam-Winget, N., Ed., Hui, J., and D. Popa, "Applicability Statement for the Routing Protocol for Low-Power and Lossy Networks (RPL) in Advanced Metering Infrastructure (AMI) Networks", RFC 8036, DOI 10.17487/RFC8036, January 2017, <<https://www.rfc-editor.org/info/rfc8036>>.
- [RFC8065] Thaler, D., "Privacy Considerations for IPv6 Adaptation-Layer Mechanisms", RFC 8065, DOI 10.17487/RFC8065, February 2017, <<https://www.rfc-editor.org/info/rfc8065>>.
- [RFC8066] Chakrabarti, S., Montenegro, G., Droms, R., and J. Woodyatt, "IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) ESC Dispatch Code Points and Guidelines", RFC 8066, DOI 10.17487/RFC8066, February 2017, <<https://www.rfc-editor.org/info/rfc8066>>.

Authors' Addresses

Jianqiang Hou
Huawei Technologies
101 Software Avenue,
Nanjing 210012
China

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

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

Email: remy.liubing@huawei.com

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

Email: yghong@etri.re.kr

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

Email: itc@sgepri.sgcc.com.cn

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

Email: charliep@computer.org

6lo
Internet-Draft
Updates: 6775 (if approved)
Intended status: Standards Track
Expires: August 27, 2018

P. Thubert, Ed.
Cisco
B. Sarikaya

M. Sethi
Ericsson
February 23, 2018

Address Protected Neighbor Discovery for Low-power and Lossy Networks
draft-ietf-6lo-ap-nd-06

Abstract

This document defines an extension to 6LoWPAN Neighbor Discovery (ND) [RFC6775][I-D.ietf-6lo-rfc6775-update] called Address Protected ND (AP-ND); AP-ND protects the owner of an address against address theft and impersonation inside a low-power and lossy network (LLN). Nodes supporting this extension compute a cryptographic Owner Unique Interface ID and associate it with one or more of their Registered Addresses. The Cryptographic ID uniquely identifies the owner of the Registered Address and can be used for proof-of-ownership. It is used in 6LoWPAN ND in place of the EUI-64-based unique ID that is associated with the registration. Once an address is registered with a Cryptographic ID, only the owner of that ID can modify the anchor state information of the Registered Address, and Source Address Validation can be enforced.

Status of This Memo

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

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

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

This Internet-Draft will expire on August 27, 2018.

Copyright Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1. Introduction	3
2. Terminology	4
3. Updating RFC 6775	5
4. New Fields and Options	5
4.1. Encoding the Public Key	5
4.2. New Crypto-ID	6
4.3. Updated EARO	6
4.4. Crypto-ID Parameters Option	8
4.5. Nonce Option	9
4.6. NDP Signature Option	9
5. Protocol Scope	9
6. Protocol Flows	10
6.1. First Exchange with a 6LR	11
6.2. Multihop Operation	13
7. Security Considerations	15
7.1. Inheriting from RTC 3971	15
7.2. Related to 6LoWPAN ND	16
7.3. OUID Collisions	16
8. IANA considerations	17
8.1. CGA Message Type	17
8.2. Crypto-Type Subregistry	17
9. Acknowledgments	17
10. References	18
10.1. Normative References	18
10.2. Informative references	19
Appendix A. Requirements Addressed in this Document	21
Authors' Addresses	21

1. Introduction

"Neighbor Discovery Optimizations for 6LoWPAN networks" [RFC6775] (6LoWPAN ND) adapts the classical IPv6 ND protocol [RFC4861][RFC4862] (IPv6 ND) for operations over a constrained low-power and lossy network (LLN). In particular, 6LoWPAN ND introduces a unicast host address registration mechanism that contributes to reduce the use of multicast messages that are present in the classical IPv6 ND protocol. 6LoWPAN ND defines a new Address Registration Option (ARO) that is carried in the unicast Neighbor Solicitation (NS) and Neighbor Advertisement (NA) messages between the 6LoWPAN Node (6LN) and the 6LoWPAN Router (6LR). Additionally, it also defines the Duplicate Address Request (DAR) and Duplicate Address Confirmation (DAC) messages between the 6LR and the 6LoWPAN Border Router (6LBR). In LLN networks, the 6LBR is the central repository of all the registered addresses in its domain.

The registration mechanism in 6LoWPAN ND [RFC6775] prevents the use of an address if that address is already present in the subnet (first come first serve). In order to validate address ownership, the registration mechanism enables the 6LR and 6LBR to validate claims for a registered address with an associated Owner Unique Interface Identifier (OUI). 6LoWPAN ND specifies that the OUI is derived from the MAC address of the device (using the 64-bit Extended Unique Identifier EUI-64 address format specified by IEEE), which can be spoofed. Therefore, any node connected to the subnet and aware of a registered-address-to-OUI mapping could effectively fake the OUI, steal the address and redirect traffic for that address towards a different 6LN. The "Update to 6LoWPAN ND" [I-D.ietf-6lo-rfc6775-update] defines an Extended ARO (EARO) option that allows to transport alternate forms of OUIs, and is a prerequisite for this specification.

According to this specification, a 6LN generates a cryptographic ID (Crypto-ID) and places it in the OUI field in the registration of one (or more) of its addresses with the 6LR(s) that the 6LN uses as default router(s). Proof of ownership of the cryptographic ID (Crypto-ID) is passed with the first registration exchange to a new 6LR, and enforced at the 6LR. The 6LR validates ownership of the cryptographic ID before it can create a registration state, or a change the anchor information, that is the Link-Layer Address and associated parameters, in an existing registration state.

The protected address registration protocol proposed in this document enables the enforcement of Source Address Validation (SAVI) [RFC7039], which ensures that only the correct owner uses a registered address in the source address field in IPv6 packets. Consequently, a 6LN that sources a packet has to use a 6LR to which

the source address of the packet is registered to forward the packet. The 6LR maintains state information for the registered address, including the MAC address, and a link-layer cryptographic key associated with the 6LN. In SAVI-enforcement mode, the 6LR allows only packets from a connected Host if the connected Host owns the registration of the source address of the packet.

The 6lo adaptation layer framework ([RFC4944], [RFC6282]) expects that a device forms its IPv6 addresses based on Layer-2 address, so as to enable a better compression. This is incompatible with "Secure Neighbor Discovery (SeND)" [RFC3971] and "Cryptographically Generated Addresses (CGAs)" [RFC3972], which derive the Interface ID (IID) in the IPv6 addresses from cryptographic material. "Privacy Considerations for IPv6 Address Generation Mechanisms" [RFC7721] places additional recommendations on the way addresses should be formed and renewed.

This document specifies that a device may form and register addresses at will, without a constraint on the way the address is formed or the number of addresses that are registered in parallel. It enables to protect multiple addresses with a single cryptographic material and to send the proof only once to a given 6LR for multiple addresses and refresher registrations.

2. Terminology

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

Readers are expected to be familiar with all the terms and concepts that are discussed in [RFC3971], [RFC3972], [RFC4861], [RFC4919], [RFC6775], and [I-D.ietf-6lo-backbone-router] which proposes an evolution of [RFC6775] for wider applicability.

This document defines Crypto-ID as an identifier of variable size which in most cases is 64 bits long. It is generated using cryptographic means explained later in this document Section 4.2. "Elliptic Curves for Security" [RFC7748] and "Edwards-Curve Digital Signature Algorithm (EdDSA)" [RFC8032] provides information on Elliptic Curve Cryptography (ECC) and a (twisted) Edwards curve, Ed25519, which can be used with this specification. "Alternative Elliptic Curve Representations" [I-D.struik-lwig-curve-representations] provides additional information on how to represent Montgomery curves and (twisted) Edwards curves as curves in short-Weierstrass form and illustrates how this can be used to implement elliptic curve computations using

existing implementations that already implement, e.g., ECDSA and ECDH using NIST [FIPS-186-4] prime curves.

The document also conforms to the terms and models described in [RFC5889] and uses the vocabulary and the concepts defined in [RFC4291] for the IPv6 Architecture. Finally, common terminology related to Low power And Lossy Networks (LLN) defined in [RFC7102] is also used.

3. Updating RFC 6775

This specification defines a cryptographic identifier (Crypto-ID) that can be used as a replacement to the MAC address in the OUID field of the EARO option; the computation of the Crypto-ID is detailed in Section 4.2. A node in possession of the necessary cryptographic material SHOULD use Crypto-ID by default as OUID in its registration. Whether a OUID is a Crypto-ID is indicated by a new "C" flag in the NS(EARO) message.

In order to prove its ownership of a Crypto-ID, the registering node needs to produce the parameters that were used to build it, as well as a nonce and a signature that will prove that it has the private key that corresponds to the public key that was used to build the Crypto-ID. This specification adds the capability to carry new options in the NS(EARO) and the NBA(EARO). These options are a variation of the CGA Option Section 4.4, a Nonce option and a variation of the RSA Signature option Section 4.6 in the NS(EARO) and a Nonce option in the NA(EARO).

4. New Fields and Options

In order to avoid an inflation of ND option types, this specification reuses / extends options defined in SEND [RFC3971] and 6LoWPAN ND [RFC6775][I-D.ietf-6lo-rfc6775-update]. This applies in particular to the CGA option and the RSA Signature Option. This specification provides aliases for the specific variations of those options as used in AP-ND. The presence of the EARO option in the NS/NA messages indicates that the options are to be understood as specified in this document. A router that would receive a NS(EARO) and try to process it as a SEND message will find that the signature does not match and drop the packet.

4.1. Encoding the Public Key

Public Key is the most important parameter in CGA Parameters (sent by 6LN in an NS message). ECC Public Key could be in uncompressed form or in compressed form where the first octet of the OCTET STRING is

0x04 and 0x02 or 0x03, respectively. Point compression can further reduce the key size by about 32 octets.

4.2. New Crypto-ID

Elliptic Curve Cryptography (ECC) is used to calculate the Crypto-ID. Each 6LN using a Crypto-ID for registration MUST have a public/private key pair. The digital signature is constructed by using the 6LN's private key over its EUI-64 (MAC) address. The signature value is computed using the ECDSA signature algorithm and the hash function used is SHA-256 [RFC6234].

NIST P-256 [FIPS186-4] that MUST be supported by all implementations. To support cryptographic algorithm agility [RFC7696], Edwards-Curve Digital Signature Algorithm (EdDSA) curve Ed25519ph (pre-hashing) [RFC8032] MAY be supported as an alternate.

The Crypto-ID is computed as follows:

1. An 8-bits modifier is selected, for instance, but not necessarily, randomly; the modifier enables a device to form multiple Crypto-IDs with a single key pair. This may be useful for privacy reasons in order to avoid the correlation of addresses based on their Crypto-ID;
2. the modifier value and the DER-encoded public key (Section 4.1) are concatenated from left to right;
3. Digital signature (SHA-256 then either NIST P-256 or EdDSA) is executed on the concatenation
4. the leftmost bits of the resulting signature are used as the Crypto-ID;

With this specification, only 64 bits are retained, but it could be expanded to more bits in the future by increasing the size of the OUID field.

4.3. Updated EARO

This specification updates the EARO option as follows:




Figure 1: Enhanced Address Registration Option

Type:	33
Length:	8-bit unsigned integer. The length of the option (including the type and length fields) in units of 8 bytes.
Status:	8-bit unsigned integer. Indicates the status of a registration in the NA response. MUST be set to 0 in NS messages. This specification uses values introduced in the update to 6LoWPAN ND [I-D.ietf-6lo-rfc6775-update], such as "Validation Requested" and "Validation Failed". No additional value is defined.
Reserved:	This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.
C:	This "C" flag is set to indicate that the Owner Unique ID field contains a Crypto-ID and that the 6LN MAY be challenged for ownership as specified in this document.
R:	Defined in [I-D.ietf-6lo-rfc6775-update].
T and TID:	Defined in [I-D.ietf-6lo-rfc6775-update].
Owner Unique ID:	When the "C" flag is set, this field contains a Crypto-ID.

4.4. Crypto-ID Parameters Option

This specification defines the Crypto-ID Parameters Option (CIPO), as a variation of the CGA Option that carries the parameters used to form a Crypto-ID. In order to provide cryptographic agility, AP-ND supports two possible signature algorithms, indicated by a Crypto-Type field. A value of 0 indicates that NIST P-256 is used for the signature operation and SHA-256 as the hash algorithm. NIST P-256 MUST be supported by all implement A value of 1 indicates that Ed25519ph is used for the signature operation and SHA-256 as the hash algorithm.




Figure 2: Crypto-ID Parameters Option

- Type: 11. This is the same value as the CGA Option, CIPO is a particular case of the CGA option
- Length: 8-bit unsigned integer. The length of the option in units of 8 octets.
- Modifier: 8-bit unsigned integer.
- Pad Length: 8-bit unsigned integer. The length of the Padding field.

Crypto-Type: The type of cryptographic algorithm used in calculation Crypto-ID. Default value of all zeros indicate NIST P-256. A value of 1 is assigned for Ed25519ph. New values may be defined later.

Public Key: Public Key of 6LN.

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

4.5. Nonce Option

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

4.6. NDP Signature Option

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

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

- o The 128-bit CGA Message Type tag [RFC3972] for AP-ND is 0x8701 55c8 0cca dd32 6ab7 e415 f148 84d0. (The tag value has been generated by the editor of this specification on random.org).
- o The signature is computed using the hash algorithm and the digital signature indicated in the Crypto-Type field of the CIP0 option using the private key associated with the public key in the CIP0.
- o The alias NDP Signature Option (NDPSO) can be used to refer to the RSAO when used as described in this specification.

5. Protocol Scope

The scope of the present work is a 6LoWPAN Low Power Lossy Network (LLN), typically a stub network connected to a larger IP network via a Border Router called a 6LBR per [RFC6775].

The 6LBR maintains a registration state for all devices in the attached LLN, and, in conjunction with the first-hop router (the 6LR), is in a position to validate uniqueness and grant ownership of an IPv6 address before it can be used in the LLN. This is a fundamental difference with a classical network that relies on IPv6

address auto-configuration [RFC4862], where there is no guarantee of ownership from the network, and any IPv6 Neighbor Discovery packet must be individually secured [RFC3971].




Figure 3: Basic Configuration

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

6. Protocol Flows

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

This specification enables to verify the ownership of the binding at any time assuming that the "C" flag is set. If it is not set, then the verification methods presented in this specification cannot be applied. The verification prevents other nodes from stealing the address and trying to attract traffic for that address or use it as their source address.

A node may use multiple IPv6 addresses at the same time. The node may use a same Crypto-ID, or multiple crypto-IDs derived from a same key pair, to protect multiple IPv6 addresses. The separation of the address and the cryptographic material avoids the constrained device

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

6.1. First Exchange with a 6LR

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

The 6LN replies to the challenge with a proof-of-ownership NS(EARO) that includes the echoed Nonce option, the CIP0 with all the parameters that were used to build EARO with a Crypto-ID, and as the last option the NDPSO with the signature. The information associated to a crypto-ID is passed to and stored by the 6LR on the first NS exchange where it appears. The 6LR SHOULD store the CIP0 information associated with the crypto-ID so it can be used for more than one address.




Figure 4: On-link Protocol Operation

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

- o Upon the first exchange with a 6LR, a 6LN may be challenged and have to produce the proof of ownership of the Crypto-ID. However, it is not expected that the proof is needed again in the periodic refresher registrations for that address, or when registering other addresses with the same OUID. When a 6LR receives a NS(EARO) registration with a new Crypto-ID as a OUID, it SHOULD challenge by responding with a NA(EARO) with a status of "Validation Requested". This process of validation MAY be skipped in networks where there is no mobility.
- o The challenge MUST also be triggered in the case of a registration for which the Source Link-Layer Address is not consistent with a state that already exists either at the 6LR or the 6LBR. In the latter case, the 6LBR returns a status of "Validation Requested" in the DAR/DAC exchange, which is echoed by the 6LR in the NA (EARO) back to the registering node. This flow should not alter a preexisting state in the 6LR or the 6LBR.

- o Upon receiving a NA(EARO) with a status of "Validation Requested", the registering node SHOULD retry its registration with a Crypto-ID Parameters Option (CIPO) Section 4.4 that contains all the necessary material for building the Crypto-ID, the Nonce and the NDP signature Section 4.6 options that prove its ownership of the Crypto-ID.
- o In order to validate the ownership, the 6LR performs the same steps as the 6LN and rebuilds the Crypto-ID based on the parameters in the CIPO. If the result is different then the validation fails. Else, the 6LR checks the signature in the NDPSO using the public key in the CIPO. If it is correct then the validation passes, else it fails.
- o If the 6LR fails to validate the signed NS(EARO), it responds with a status of "Validation Failed". After receiving a NA(EARO) with a status of "Validation Failed", the registering node SHOULD try an alternate Signature Algorithm and Crypto-ID. In any case, it MUST NOT use this Crypto-ID for registering with that 6LR again.

6.2. Multihop Operation

In a multihop 6LoWPAN, the registration with Crypto-ID is propagated to 6LBR as described in Section 6.2. If a chain of trust is present between the 6LR and the 6LBR, then there is no need to propagate the proof of ownership to the 6LBR. All the 6LBR needs to know is that this particular OUID is randomly generated, so as to enforce that any update via a different 6LR is also random.

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

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

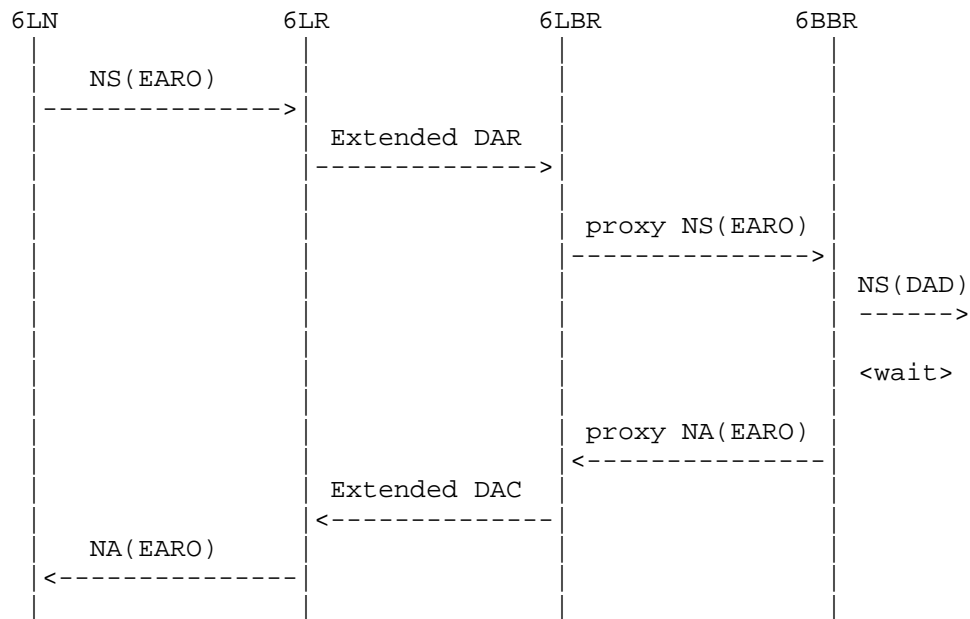


Figure 5: (Re-)Registration Flow

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

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

Occasionally, a 6LR might miss the node's OUID (that it received in ARO). 6LR should be able to ask for it again. This is done by restarting the exchanges shown in Figure 4. The result enables 6LR to refresh the information that was lost. The 6LR MUST send DAR message with ARO to 6LBR. The 6LBR replies with a DAC message with the information copied from the DAR, and the Status field is set to zero. With this exchange, the 6LBR can (re)validate and store the information to make sure that the 6LR is not a fake.

In some cases, the 6LBR may use a DAC message to solicit a Crypto-ID from a 6LR and also requests 6LR to verify the EUI-64 6LR received from 6LN. This may happen when a 6LN node is compromised and a fake node is sending the Crypto-ID as if it is the node's EUI-64. Note that the detection in this case can only be done by 6LBR not by 6LR.

7. Security Considerations

7.1. Inheriting from RTC 3971

The observations regarding the threats to the local network in [RFC3971] also apply to this specification. Considering RFC3971 security section subsection by subsection:

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

Neighbor Unreachability Detection Failure With RFC6775, a NUD can still be used by the endpoint to assess the liveness of a device. The NUD request may be protected by SEND in which case the provision in section 9.2.2. of RFC 3972 applies. The response to the NUD may be proxied by a backbone router only if it has a fresh registration state for it. The registration being protected by this specification, the proxied NUD response provides a truthful information on the original owner of the address but it cannot be proven using SEND. If the NUD response is not proxied, the 6LR will pass the lookup to the end device which will respond with a traditional NA. If the 6LR does not have a cache entry associated for the device, it can issue a NA with EARO (status=Validation Requested) upon the NA from the device, which will trigger a NS that will recreate and revalidate the ND cache entry.

Duplicate Address Detection DoS Attack Inside the LLN, Duplicate Addresses are sorted out using the OUID, which differentiates it from a movement. DAD coming from the backbone are not forwarded over the LLN so the LLN is protected by the backbone routers. Over the backbone, the EARO option is present in NS/NA messages. This protects against misinterpreting a movement for a duplication, and enables to decide which backbone router has the freshest registration and thus most possibly the device attached to it. But this specification does not guarantee that the backbone router claiming an address over the backbone is not an attacker.

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

Replay Attacks A Nonce given by the 6LR in the NA with EARO (status=Validation Requested) and echoed in the signed NS guarantees against replay attacks of the NS(EARO). The NA(EARO) is not protected and can be forged by a rogue node that is not the 6LR in order to force the 6LN to rebuild a NS(EARO) with the proof of ownership, but that rogue node must have access to the L2 radio network next to the 6LN to perform the attack.

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

7.2. Related to 6LoWPAN ND

The threats discussed in 6LoWPAN ND [RFC6775] and its update [I-D.ietf-6lo-rfc6775-update] also apply here. Compared with SEND, this specification saves about 1Kbyte in every NS/NA message. Also, this specification separates the cryptographic identifier from the registered IPv6 address so that a node can have more than one IPv6 address protected by the same cryptographic identifier. SEND forces the IPv6 address to be cryptographic since it integrates the CGA as the IID in the IPv6 address. This specification frees the device to form its addresses in any fashion, so as to enable the classical 6LoWPAN compression which derives IPv6 addresses from Layer-2 addresses, as well as privacy addresses. The threats discussed in Section 9.2 of [RFC3971] are countered by the protocol described in this document as well.

7.3. OUID Collisions

Collisions of Owner Unique Interface Identifier (OUID) (which is the Crypto-ID in this specification) is a possibility that needs to be considered. The formula for calculating the probability of a collision is $1 - e^{-k^2/(2n)}$ where n is the maximum population size (2^{64} here, $1.84E19$) and K is the actual population (number of nodes). If the Crypto-ID is 64-bit long, then the chance of finding a collision is 0.01% when the network contains 66 million nodes. It is important to note that the collision is only relevant when this happens within one stub network (6LBR). A collision of Crypto-ID is

a rare event. In the case of a collision, an attacker may be able to claim the registered address of an another legitimate node. However for this to happen, the attacker would also need to know the address which was registered by the legitimate node. This registered address is however never broadcasted on the network and therefore it provides an additional entropy of 64-bits that an attacker must correctly guess. To prevent such a scenario, it is RECOMMENDED that nodes derive the address being registered independently of the OUID.

8. IANA considerations

8.1. CGA Message Type

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

8.2. Crypto-Type Subregistry

IANA is requested to create a new subregistry "Crypto-Type Subregistry" in the "Internet Control Message Protocol version 6 (ICMPv6) Parameters". The registry is indexed by an integer 0..255 and contains a Signature Algorithm and a Hash Function as shown in Table 1. The following Crypto-Type values are defined in this document:

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

Table 1: Crypto-Types

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

9. Acknowledgments

Many thanks to Charlie Perkins for his in-depth review and constructive suggestions. We are also especially grateful to Rene Struik and Robert Moskowitz for their comments that lead to many improvements to this document, in particular WRT ECC computation and references.

10. References

10.1. Normative References

- [FIPS-186-4]
FIPS 186-4, "Digital Signature Standard (DSS), Federal Information Processing Standards Publication 186-4", US Department of Commerce/National Institute of Standards and Technology Gaithersburg, MD, July 2013.
- [I-D.ietf-6lo-rfc6775-update]
Thubert, P., Nordmark, E., Chakrabarti, S., and C. Perkins, "An Update to 6LoWPAN ND", draft-ietf-6lo-rfc6775-update-13 (work in progress), February 2018.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3279] Bassham, L., Polk, W., and R. Housley, "Algorithms and Identifiers for the Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 3279, DOI 10.17487/RFC3279, April 2002, <<https://www.rfc-editor.org/info/rfc3279>>.
- [RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "Secure Neighbor Discovery (SEND)", RFC 3971, DOI 10.17487/RFC3971, March 2005, <<https://www.rfc-editor.org/info/rfc3971>>.
- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, DOI 10.17487/RFC3972, March 2005, <<https://www.rfc-editor.org/info/rfc3972>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.
- [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>>.

- [RFC5758] Dang, Q., Santesson, S., Moriarty, K., Brown, D., and T. Polk, "Internet X.509 Public Key Infrastructure: Additional Algorithms and Identifiers for DSA and ECDSA", RFC 5758, DOI 10.17487/RFC5758, January 2010, <<https://www.rfc-editor.org/info/rfc5758>>.
- [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>>.

10.2. Informative references

- [FIPS186-4] "FIPS Publication 186-4: Digital Signature Standard", July 2013, <<http://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.186-4.pdf>>.
- [I-D.ietf-6lo-backbone-router] Thubert, P., "IPv6 Backbone Router", draft-ietf-6lo-backbone-router-05 (work in progress), January 2018.
- [I-D.struik-lwig-curve-representations] Struik, R., "Alternative Elliptic Curve Representations", draft-struik-lwig-curve-representations-00 (work in progress), November 2017.
- [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>>.
- [RFC5889] Baccelli, E., Ed. and M. Townsley, Ed., "IP Addressing Model in Ad Hoc Networks", RFC 5889, DOI 10.17487/RFC5889, September 2010, <<https://www.rfc-editor.org/info/rfc5889>>.
- [RFC6234] Eastlake 3rd, D. and T. Hansen, "US Secure Hash Algorithms (SHA and SHA-based HMAC and HKDF)", RFC 6234, DOI 10.17487/RFC6234, May 2011, <<https://www.rfc-editor.org/info/rfc6234>>.

- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<https://www.rfc-editor.org/info/rfc6282>>.
- [RFC7039] Wu, J., Bi, J., Bagnulo, M., Baker, F., and C. Vogt, Ed., "Source Address Validation Improvement (SAVI) Framework", RFC 7039, DOI 10.17487/RFC7039, October 2013, <<https://www.rfc-editor.org/info/rfc7039>>.
- [RFC7102] Vasseur, JP., "Terms Used in Routing for Low-Power and Lossy Networks", RFC 7102, DOI 10.17487/RFC7102, January 2014, <<https://www.rfc-editor.org/info/rfc7102>>.
- [RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <<https://www.rfc-editor.org/info/rfc7217>>.
- [RFC7696] Housley, R., "Guidelines for Cryptographic Algorithm Agility and Selecting Mandatory-to-Implement Algorithms", BCP 201, RFC 7696, DOI 10.17487/RFC7696, November 2015, <<https://www.rfc-editor.org/info/rfc7696>>.
- [RFC7721] Cooper, A., Gont, F., and D. Thaler, "Security and Privacy Considerations for IPv6 Address Generation Mechanisms", RFC 7721, DOI 10.17487/RFC7721, March 2016, <<https://www.rfc-editor.org/info/rfc7721>>.
- [RFC7748] Langley, A., Hamburg, M., and S. Turner, "Elliptic Curves for Security", RFC 7748, DOI 10.17487/RFC7748, January 2016, <<https://www.rfc-editor.org/info/rfc7748>>.
- [RFC8032] Josefsson, S. and I. Liusvaara, "Edwards-Curve Digital Signature Algorithm (EdDSA)", RFC 8032, DOI 10.17487/RFC8032, January 2017, <<https://www.rfc-editor.org/info/rfc8032>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.

Appendix A. Requirements Addressed in this Document

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

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

Authors' Addresses

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

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

Behcet Sarikaya
Plano, TX
USA

Email: sarikaya@ieee.org

Mohit Sethi
Ericsson
Hirsalantie
Jorvas 02420

Email: mohit@piuha.net

6lo
Internet-Draft
Intended status: Standards Track
Expires: August 27, 2018

P. Thubert, Ed.
cisco
February 23, 2018

IPv6 Backbone Router
draft-ietf-6lo-backbone-router-06

Abstract

This specification proposes proxy operations for IPv6 Neighbor Discovery on behalf of devices located on broadcast-inefficient wireless networks. A broadcast-efficient backbone running classical IPv6 Neighbor Discovery federates multiple wireless links to form a large MultiLink Subnet, but the broadcast domain does not need to extend to the wireless links for the purpose of ND operation. Backbone Routers placed at the wireless edge of the backbone proxy the ND operation and route packets from/to registered nodes, and wireless nodes register or are proxy-registered to the Backbone Router to setup proxy services in a fashion that is essentially similar to a classical Layer-2 association.

Status of This Memo

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

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

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

This Internet-Draft will expire on August 27, 2018.

Copyright Notice

Copyright (c) 2018 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. Applicability and Requirements Served	4
3. Terminology	6
4. Overview	7
5. Backbone Router Routing Operations	9
5.1. Over the Backbone Link	10
5.2. Over the LLN Link	11
6. BackBone Router Proxy Operations	13
6.1. Registration and Binding State Creation	15
6.2. Defending Addresses	17
7. Security Considerations	18
8. Protocol Constants	18
9. IANA Considerations	18
10. Acknowledgments	19
11. References	19
11.1. Normative References	19
11.2. Informative References	20
11.3. External Informative References	23
Appendix A. Requirements	24
A.1. Requirements Related to Mobility	24
A.2. Requirements Related to Routing Protocols	25
A.3. Requirements Related to the Variety of Low-Power Link types	26
A.4. Requirements Related to Proxy Operations	26
A.5. Requirements Related to Security	27
A.6. Requirements Related to Scalability	28
Author's Address	29

1. Introduction

One of the key services provided by IEEE std. 802.1 [IEEEstd8021] Ethernet Bridging is an efficient and reliable broadcast service, and multiple applications and protocols have been built that heavily depend on that feature for their core operation. But a wide range of wireless networks do not provide the solid and cheap broadcast capabilities of Ethernet Bridging, and protocols designed for bridged networks that rely on broadcast often exhibit disappointing behaviours when applied unmodified to a wireless medium.

IEEE std. 802.11 [IEEEstd80211] Access Points (APs) deployed in an Extended Service Set (ESS) effectively act as bridges, but, in order to ensure a solid connectivity to the devices and protect the medium against harmful broadcasts, they refrain from relying on broadcast-intensive protocols such as Transparent Bridging on the wireless side. Instead, an association process is used to register proactively the MAC addresses of the wireless device (STA) to the AP, and then the APs proxy the bridging operation and cancel the broadcasts.

Classical IPv6 [RFC8200] Neighbor Discovery [RFC4861] [RFC4862] Protocol (NDP) operations are reactive and rely heavily on multicast operations to locate an on-link correspondent and ensure address uniqueness, which is a pillar that sustains the whole IP architecture. When the Duplicate Address Detection [RFC4862] (DAD) mechanism was designed, it was a natural match with the efficient broadcast operation of Ethernet Bridging, but with the unreliable broadcast that is typical of wireless media, DAD is bound to fail to discover duplications [I-D.yourtchenko-6man-dad-issues]. In other words, because the broadcast service is unreliable, DAD appears to work on wireless media not because address duplication is detected and solved as designed, but because the duplication is a very rare event as a side effect of the sheer amount of entropy in 64-bits Interface IDs.

In the real world, IPv6 multicast messages are effectively broadcast, so they are processed by most if not all wireless nodes over the ESS fabric even when very few if any of the nodes is effectively listening to the multicast address. It results that a simple Neighbor Solicitation (NS) lookup message [RFC4861], that is supposedly targeted to a very small group of nodes, ends up polluting the whole wireless bandwidth across the fabric [I-D.vyncke-6man-mcast-not-efficient]. In other words, the reactive IPv6 ND operation leads to undesirable power consumption in battery-operated devices.

The inefficiencies of using radio broadcasts to support IPv6 NDP lead the community to consider (again) splitting the broadcast domain between the wired and the wireless access links. One classical way to achieve this is to split the subnet in multiple ones, and at the extreme provide a /64 per wireless device. Another is to proxy the Layer-3 protocols that rely on broadcast operation at the boundary of the wired and wireless domains, effectively emulating the Layer-2 association at layer-3. To that effect, the current IEEE std. 802.11 specifications require the capability to perform ARP and ND proxy [RFC4389] functions at the Access Points (APs).

But for the lack a comprehensive specification for the ND proxy and in particular the lack of an equivalent to an association process, implementations have to rely on snooping for acquiring the related state, which is unsatisfactory in a lossy and mobile conditions. With snooping, a state (e.g. a new IPv6 address) may not be discovered or a change of state (e.g. a movement) may be missed, leading to unreliable connectivity.

In the context of IEEE std. 802.15.4 [IEEEstd802154], the step of considering the radio as a medium that is different from Ethernet was already taken with the publication of Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs) [RFC6775]. RFC 6775 is updated as [I-D.ietf-6lo-rfc6775-update]; the update includes changes that are required by this document.

This specification applies that same thinking to other wireless links such as Low-Power IEEE std. 802.11 (Wi-Fi) and IEEE std. 802.15.1 (Bluetooth) [IEEEstd802151], and extends [RFC6775] to enable proxy operation by the 6BBR so as to decouple the broadcast domain in the backbone from the wireless links. The proxy operation can be maintained asynchronous so that low-power nodes or nodes that are deep in a mesh do not need to be bothered synchronously when a lookup is performed for their addresses, effectively implementing the ND contribution to the concept of a Sleep Proxy [I-D.nordmark-6man-dad-approaches].

2. Applicability and Requirements Served

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

This specification updates and generalizes 6LoWPAN ND to a broader range of Low power and Lossy Networks (LLNs) with a solid support for Duplicate Address Detection (DAD) and address lookup that does not require broadcasts over the LLNs. The term LLN is used loosely in this specification to cover multiple types of WLANs and WPANs, including Low-Power Wi-Fi, BLUETOOTH(R) Low Energy, IEEE std. 802.11AH and IEEE std. 802.15.4 wireless meshes, so as to address the requirements listed in Appendix A.3

The scope of this draft is a Backbone Link that federates multiple LLNs as a single IPv6 MultiLink Subnet. Each LLN in the subnet is anchored at an IPv6 Backbone Router (6BBR). The Backbone Routers interconnect the LLNs over the Backbone Link and emulate that the LLN nodes are present on the Backbone using proxy-ND operations. This specification extends IPv6 ND over the backbone to discriminate address movement from duplication and eliminate stale state in the backbone routers and backbone nodes once a LLN node has roamed. This way, mobile nodes may roam rapidly from a 6BBR to the next and requirements in Appendix A.1 are met.

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

The Link Layer Address (LLA) that is returned as Target LLA (TLA) in Neighbor Advertisements (NA) messages by the 6BBR on behalf of the Registered Node over the backbone may be that of the Registering Node, in which case the 6BBR needs to bridge the unicast packets (Bridging proxy), or that of the 6BBR on the backbone, in which case the 6BBR needs to route the unicast packets (Routing proxy). In the latter case, the 6BBR may maintain the list of correspondents to which it has advertised its own MAC address on behalf of the LLN node and the IPv6 ND operation is minimized as the number of nodes scale up in the LLN. This enables to meet the requirements in Appendix A.6 as long as the 6BBRs are dimensioned for the number of registration that each needs to support.

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

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

BTLE, the 6BBR is collocated with the router that implements the BTLE central role as discussed in section 2.2 of [RFC7668].

3. Terminology

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

Readers are expected to be familiar with all the terms and concepts that are discussed in "Neighbor Discovery for IP version 6" [RFC4861], "IPv6 Stateless Address Autoconfiguration" [RFC4862], "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals" [RFC4919], "Neighbor Discovery Optimization for Low-power and Lossy Networks" [RFC6775] and "Multi-link Subnet Support in IPv6" [I-D.ietf-ipv6-multilink-subnets].

Readers would benefit from reading "Multi-Link Subnet Issues" [RFC4903], "Mobility Support in IPv6" [RFC6275], "Neighbor Discovery Proxies (ND Proxy)" [RFC4389] and "Optimistic Duplicate Address Detection" [RFC4429] prior to this specification for a clear understanding of the art in ND-proxying and binding.

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

Sleeping Proxy A 6BBR acts as a Sleeping Proxy if it answers ND Neighbor Solicitation over the backbone on behalf of the Registered Node whenever possible. This is the default mode for this specification but it may be overridden, for instance by configuration, into Unicasting Proxy.

Unicasting Proxy As a Unicasting Proxy, the 6BBR forwards NS messages to the Registering Node, transforming Layer-2 multicast into unicast whenever possible.

Routing proxy A 6BBR acts as a routing proxy if it advertises its own MAC address, as opposed to that of the node that performs the registration, as the TLLA in the proxied NAs over the backbone. In that case, the MAC address of the node is not visible at Layer-2 over the backbone and the bridging fabric is not aware of the addresses of the LLN devices and their mobility. The 6BBR installs a connected host route towards the registered node over the interface to the node, and acts as a Layer-3 router for unicast packets to the node. The 6BBR updates the ND Neighbor Cache Entries (NCE) in correspondent

nodes if the wireless node moves and registers to another 6BBR, either with a single broadcast, or with a series of unicast NA(O) messages, indicating the TLLA of the new router.

Bridging proxy A 6BBR acts as a bridging proxy if it advertises the MAC address of the node that performs the registration as the TLLA in the proxied NAs over the backbone. In that case, the MAC address and the mobility of the node is still visible across the bridged backbone fabric, as is traditionally the case with Layer-2 APs. The 6BBR acts as a Layer-2 bridge for unicast packets to the registered node. The MAC address exposed in the S/TLLA is that of the Registering Node, which is not necessarily the Registered Device. When a device moves within a LLN mesh, it may end up attached to a different 6LBR acting as Registering Node, and the LLA that is exposed over the backbone will change.

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

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

4. Overview

An LLN node can move freely from an LLN anchored at a Backbone Router to an LLN anchored at another Backbone Router on the same backbone and conserve any of the IPv6 addresses that it has formed, transparently.




Figure 1: Backbone Link and Backbone Routers

The Backbone Routers maintain an abstract Binding Table of their Registered Nodes. The Binding Table operates as a distributed database of all the wireless Nodes whether they reside on the LLNs or on the backbone, and use an extension to the Neighbor Discovery Protocol to exchange that information across the Backbone in the classical ND reactive fashion.

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

Address duplication is sorted out with the Owner Unique-ID field in the EARO, which is a generalization of the EUI-64 that allows different types of unique IDs beyond the name space derived from the MAC addresses. First-Come First-Serve rules apply, whether the duplication happens between LLN nodes as represented by their respective 6BBRs, or between an LLN node and a classical node that defends its address over the backbone with classical ND and does not include the EARO option.

In case of conflicting registrations to multiple 6BBRs from a same node, a sequence counter called Transaction ID (TID) in the EARO enables 6BBRs to sort out the latest anchor for that node.

Registrations with a same TID are compatible and maintained, but, in case of different TIDs, only the freshest registration is maintained and the stale state is eliminated. The EARO also transports a 'R' flag to be used by a 6LN when registering, to indicate that this 6LN is not a router and that it will not handle its own reachability.

With this specification, Backbone Routers perform a ND proxy operation over the Backbone Link on behalf of their Registered Nodes. The registration to the proxy service is done with a NS/NA(EARO) exchange. The EARO option with a 'R' flag is used in this specification to indicate to the 6BBR that it is expected to perform this proxy operation. The Backbone Router operation is essentially similar to that of a Mobile IPv6 (MIPv6) [RFC6275] Home Agent. This enables mobility support for LLN nodes that would move outside of the network delimited by the Backbone link attach to a Home Agent from that point on. This also enables collocation of Home Agent functionality within Backbone Router functionality on the same backbone interface of a router. Further specification may extend this by allowing the 6BBR to redistribute host routes in routing protocols that would operate over the backbone, or in MIPv6 or the Locator/ID Separation Protocol (LISP) [RFC6830] to support mobility on behalf of the nodes, etc...

The Optimistic Duplicate Address Detection [RFC4429] (ODAD) specification details how an address can be used before a Duplicate Address Detection (DAD) is complete, and insists that an address that is TENTATIVE should not be associated to a Source Link-Layer Address Option in a Neighbor Solicitation message. This specification leverages ODAD to create a temporary proxy state in the 6BBR till DAD is completed over the backbone. This way, the specification enables to distribute proxy states across multiple 6BBR and co-exist with classical ND over the backbone.

5. Backbone Router Routing Operations




Figure 2: Routing Configuration in the ML Subnet

5.1. Over the Backbone Link

The Backbone Router is a specific kind of Border Router that performs proxy Neighbor Discovery on its backbone interface on behalf of the nodes that it has discovered on its LLN interfaces.

The backbone is expected to be a high speed, reliable Backbone link, with affordable and reliable multicast capabilities, such as a bridged Ethernet Network, and to allow a full support of classical ND as specified in [RFC4861] and subsequent RFCs. In other words, the backbone is not a LLN.

Still, some restrictions of the attached LLNs will apply to the backbone. In particular, it is expected that the MTU is set to the same value on the backbone and all attached LLNs, and the scalability of the whole subnet requires that broadcast operations are avoided as much as possible on the backbone as well. Unless configured otherwise, the Backbone Router MUST echo the MTU that it learns in RAs over the backbone in the RAs that it sends towards the LLN links.

As a router, the Backbone Router behaves like any other IPv6 router on the backbone side. It has a connected route installed towards the backbone for the prefixes that are present on that backbone and that it proxies for on the LLN interfaces.

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

When possible, the 6BBR SHOULD use unicast or solicited-node multicast address (SNMA) [RFC4291] to defend its Registered Addresses over the backbone. In particular, the 6BBR MUST join the SNMA group that corresponds to a Registered Address as soon as it creates an entry for that address and as long as it maintains that entry, whatever the state of the entry. The expectation is that it is possible to get a message delivered to all the nodes on the backbone that listen to a particular address and support this specification - which includes all the 6BBRs in the MultiLink Subnet - by sending a multicast message to the associated SNMA over the backbone.

The support of Optimistic DAD (ODAD) [RFC4429] is recommended for all nodes in the backbone and followed by the 6BBRs in their proxy activity over the backbone. With ODAD, any optimistic node MUST join the SNMA of a Tentative address, which interacts better with this specification.

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

Since the MultiLink Subnet may grow very large in terms of individual IPv6 addresses, multicasts should be avoided as much as possible even on the backbone. Though it is possible for plain hosts to participate with legacy IPv6 ND support, the support by all nodes connected to the backbone of [I-D.ietf-6man-rs-refresh] is recommended, and this implies the support of [RFC7559] as well.

5.2. Over the LLN Link

As a router, the Nodes and Backbone Router operation on the LLN follows [RFC6775]. Per that specification, LLN Hosts generally do not depend on multicast RAs to discover routers. It is still generally required for LLN nodes to accept multicast RAs [RFC7772], but those are rare on the LLN link. Nodes are expected to follow the

Simple Procedures for Detecting Network Attachment in IPv6 [RFC6059] (DNA procedures) to assert movements, and to support the Packet-Loss Resiliency for Router Solicitations [RFC7559] to make the unicast RS more reliable.

An LLN node signals that it requires IPv6 ND proxy services from a 6BBR by registering the corresponding IPv6 Address with an NS(EARO) message with the 'R' flag set. The LLN node that performs the registration (the Registering Node) may be the owner of the IPv6 Address (the Registered Node) or a 6LBR that performs the registration on its behalf.

When operating as a Routing Proxy, the router installs hosts routes (/128) to the Registered Addresses over the LLN links, via the Registering Node as identified by the Source Address and the SLLAO option in the NS(EARO) messages.

In that mode, the 6BBR handles the ND protocol over the backbone on behalf of the Registered Nodes, using its own MAC address in the TLLA and SLLA options in proxied NS and NA messages. It results that for each Registered Address, a number of peer Nodes on the backbone have resolved the address with the 6BBR MAC address and keep that mapping stored in their Neighbor cache.

The 6BBR SHOULD maintain, per Registered Address, the list of the peers on the backbone to which it answered with its MAC address, and when a binding moves to a different 6BBR, it SHOULD send a unicast gratuitous NA(O) individually to each of them to inform them that the address has moved and pass the MAC address of the new 6BBR in the TLLAO option. If the 6BBR can not maintain that list, then it SHOULD remember whether that list is empty or not and if not, send a multicast NA(O) to all nodes to update the impacted Neighbor Caches with the information from the new 6BBR.

The Bridging Proxy is a variation where the BBR function is implemented in a Layer-3 switch or an wireless Access Point that acts as a Host from the IPv6 standpoint, and, in particular, does not operate the routing of IPv6 packets. In that case, the SLLAO in the proxied NA messages is that of the Registering Node and classical bridging operations take place on data frames.

If a registration moves from one 6BBR to the next, but the Registering Node does not change, as indicated by the S/TLLAO option in the ND exchanges, there is no need to update the Neighbor Caches in the peers Nodes on the backbone. On the other hand, if the LLAO changes, the 6BBR SHOULD inform all the relevant peers as described above, to update the impacted Neighbor Caches. In the same fashion,

if the Registering Node changes with a new registration, the 6BBR SHOULD also update the impacted Neighbor Caches over the backbone.

6. BackBone Router Proxy Operations

This specification enables a Backbone Router to proxy Neighbor Discovery operations over the backbone on behalf of the nodes that are registered to it, allowing any node on the backbone to reach a Registered Node as if it was on-link. The backbone and the LLNs are considered different Links in a MultiLink subnet but the prefix that is used may still be advertised as on-link on the backbone to support legacy nodes; multicast ND messages are link-scoped and not forwarded across the backbone routers.

ND Messages on the backbone side that do not match to a registration on the LLN side are not acted upon on the LLN side, which stands protected. On the LLN side, the prefixes associated to the MultiLink Subnet are presented as not on-link, so address resolution for other hosts do not occur.

The default operation in this specification is Sleeping proxy which means:

- o creating a new entry in an abstract Binding Table for a new Registered Address and validating that the address is not a duplicate over the backbone
- o defending a Registered Address over the backbone using NA messages with the Override bit set on behalf of the sleeping node whenever possible
- o advertising a Registered Address over the backbone using NA messages, asynchronously or as a response to a Neighbor Solicitation messages.
- o Looking up a destination over the backbone in order to deliver packets arriving from the LLN using Neighbor Solicitation messages.
- o Forwarding packets from the LLN over the backbone, and the other way around.
- o Eventually triggering a liveness verification of a stale registration.

A 6BBR may act as a Sleeping Proxy only if the state of the binding entry is REACHABLE, or TENTATIVE in which case the answer is delayed.

In any other state, the Sleeping Proxy operates as a Unicasting Proxy.

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

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

The Backbone Routers maintain a distributed binding table, using classical ND over the backbone to detect duplication. This specification requires that:

1. All addresses that can be reachable from the backbone, including IPv6 addresses based on burn-in EUI64 addresses MUST be registered to the 6BBR.
2. A Registered Node MUST include the EARO option in an NS message that used to register an addresses to a 6LR; the 6LR MUST propagate that option unchanged to the 6LBR in the DAR/DAC exchange, and the 6LBR MUST propagate that option unchanged in proxy registrations.
3. The 6LR MUST echo the same EARO option in the NA that it uses to respond, but for the status filed which is not used in NS messages, and significant in NA.

A false positive duplicate detection may arise over the backbone, for instance if the Registered Address is registered to more than one LBR, or if the node has moved. Both situations are handled gracefully unbeknownst to the node. In the former case, one LBR

becomes primary to defend the address over the backbone while the others become secondary and may still forward packets back and forth. In the latter case the LBR that receives the newest registration wins and becomes primary.

The expectation in this specification is that there is a single Registering Node at a time per Backbone Router for a given Registered Address, but that a Registered Address may be registered to Multiple 6BBRs for higher availability.

Over the LLN, and for any given Registered Address, it is REQUIRED that:

- de-registrations (newer TID, same OUID, null Lifetime) are accepted and responded immediately with a status of 4; the entry is deleted;

- newer registrations (newer TID, same OUID, non-null Lifetime) are accepted and responded with a status of 0 (success); the entry is updated with the new TID, the new Registration Lifetime and the new Registering Node, if any has changed; in TENTATIVE state the response is held and may be overwritten; in other states the Registration-Lifetime timer is restarted and the entry is placed in REACHABLE state.

- identical registrations (same TID, same OUID) from a same Registering Node are not processed but responded with a status of 0 (success); they are expected to be identical and an error may be logged if not; in TENTATIVE state, the response is held and may be overwritten, but it MUST be eventually produced and it carries the result of the DAD process;

- older registrations (not(newer or equal) TID, same OUID) from a same Registering Node are ignored;

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

- and any registration for a different Registered Node (different OUID) are responded immediately with a status of 1 (duplicate).

6.1. Registration and Binding State Creation

Upon a registration for a new address with an NS(EARO) with the 'R' bit set, the 6BBR performs a DAD operation over the backbone placing the new address as target in the NS-DAD message. The EARO from the

registration MUST be placed unchanged in the NS-DAD message, and an entry is created in TENTATIVE state for a duration of `TENTATIVE_DURATION`. The NS-DAD message is sent multicast over the backbone to the SNMA address associated with the registered address. If that operation is known to be costly, and the 6BBR has an indication from another source (such as a NCE) that the Registered Address was present on the backbone, that information may be leveraged to send the NS-DAD message as a Layer-2 unicast to the MAC that was associated with the Registered Address.

In TENTATIVE state:

- o the entry is removed if an NA is received over the backbone for the Registered Address with no EARO option, or with an EARO option with a status of 1 (duplicate) that indicates an existing registration for another LLN node. The OUID and TID fields in the EARO option received over the backbone are ignored. A status of 1 is returned in the EARO option of the NA back to the Registering Node;
- o the entry is also removed if an NA with an ARO option with a status of 3 (moved), or a NS-DAD with an ARO option that indicates a newer registration for the same Registered Node, is received over the backbone for the Registered Address. A status of 3 is returned in the NA(EARO) back to the Registering Node;
- o when a registration is updated but not deleted, e.g. from a newer registration, the DAD process on the backbone continues and the running timers are not restarted;
- o Other NS (including DAD with no EARO option) and NA from the backbone are not responded in TENTATIVE state, but the list of their origins may be kept in memory and if so, the 6BBR may send them each a unicast NA with eventually an EARO option when the `TENTATIVE_DURATION` timer elapses, so as to cover legacy nodes that do not support ODAD.
- o When the `TENTATIVE_DURATION` timer elapses, a status 0 (success) is returned in a NA(EARO) back to the Registering Node(s), and the entry goes to REACHABLE state for the Registration Lifetime; the DAD process is successful and the 6BBR MUST send a multicast NA(EARO) to the SNMA associated to the Registered Address over the backbone with the Override bit set so as to take over the binding from other 6BBRs.

6.2. Defending Addresses

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

- o If the 6BBR is primary, or does not support the function of primary, it MUST defend that address over the backbone upon an incoming NS-DAD, either if the NS does not carry an EARO, or if an EARO is present that indicates a different Registering Node (different OUID). The 6BBR sends a NA message with the Override bit set and the NA carries an EARO option if and only if the NS-DAD did so. When present, the EARO in the NA(O) that is sent in response to the NS-DAD(EARO) carries a status of 1 (duplicate), and the OUID and TID fields in the EARO option are obfuscated with null or random values to avoid network scanning and impersonation attacks.
- o If the 6BBR receives an NS-DAD(EARO) that reflect a newer registration, the 6BBR updates the entry and the routing state to forward packets to the new 6BBR, but keeps the entry REACHABLE. In that phase, it MAY use REDIRECT messages to reroute traffic for the Registered Address to the new 6BBR.
- o If the 6BBR receives an NA(EARO) that reflect a newer registration, the 6BBR removes its entry and sends a NA(AERO) with a status of 3 (moved) to the Registering Node, if the Registering Node is different from the Registered Node. If necessary, the 6BBR cleans up ND cache in peers nodes as discussed in Section 5.1, by sending a series of unicast to the impacted nodes, or one broadcast NA(O) to all-nodes.
- o If the 6BBR received a NS(LOOKUP) for a Registered Address, it answers immediately with an NA on behalf of the Registered Node, without polling it. There is no need of an EARO in that exchange.
- o When the Registration-Lifetime timer elapses, the entry goes to STALE state for a duration of STABLE_STALE_DURATION in LLNs that keep stable addresses such as LWPANs, and UNSTABLE_STALE_DURATION in LLNs where addresses are renewed rapidly, e.g. for privacy reasons.

The STALE state is a chance to keep track of the backbone peers that may have an ND cache pointing on this 6BBR in case the Registered Address shows back up on this or a different 6BBR at a later time. In STALE state:

- o If the Registered Address is claimed by another node on the backbone, with an NS-DAD or an NA, the 6BBR does not defend the address. Upon an NA(O), or the stale time elapses, the 6BBR

removes its entry and sends a NA(AERO) with a status of 4 (removed) to the Registering Node.

- o If the 6BBR received a NS(LOOKUP) for a Registered Address, the 6BBR MUST send an NS(NUD) following rules in [RFC7048] to the Registering Node targeting the Registered Address prior to answering. If the NUD succeeds, the operation in REACHABLE state applies. If the NUD fails, the 6BBR refrains from answering the lookup. The NUD expected to be mapped by the Registering Node into a liveliness validation of the Registered Node if they are in fact different nodes.

7. Security Considerations

This specification expects that the link layer is sufficiently protected, either by means of physical or IP security for the Backbone Link or MAC sublayer cryptography. In particular, it is expected that the LLN MAC provides secure unicast to/from the Backbone Router and secure Broadcast from the Backbone Router in a way that prevents tempering with or replaying the RA messages.

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

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

8. Protocol Constants

This Specification uses the following constants:

TENTATIVE_DURATION:	800 milliseconds
STABLE_STALE_DURATION:	24 hours
UNSTABLE_STALE_DURATION:	5 minutes
DEFAULT_NS_POLLING:	3 times

9. IANA Considerations

This document has no request to IANA.

10. Acknowledgments

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

11. References

11.1. Normative References

- [I-D.ietf-6lo-rfc6775-update]
Thubert, P., Nordmark, E., Chakrabarti, S., and C. Perkins, "An Update to 6LoWPAN ND", draft-ietf-6lo-rfc6775-update-13 (work in progress), February 2018.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.
- [RFC4429] Moore, N., "Optimistic Duplicate Address Detection (DAD) for IPv6", RFC 4429, DOI 10.17487/RFC4429, April 2006, <<https://www.rfc-editor.org/info/rfc4429>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <<https://www.rfc-editor.org/info/rfc4862>>.
- [RFC6059] Krishnan, S. and G. Daley, "Simple Procedures for Detecting Network Attachment in IPv6", RFC 6059, DOI 10.17487/RFC6059, November 2010, <<https://www.rfc-editor.org/info/rfc6059>>.
- [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>>.
- [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>>.

11.2. Informative References

- [I-D.chakrabarti-nordmark-6man-efficient-nd]
Chakrabarti, S., Nordmark, E., Thubert, P., and M. Wasserman, "IPv6 Neighbor Discovery Optimizations for Wired and Wireless Networks", draft-chakrabarti-nordmark-6man-efficient-nd-07 (work in progress), February 2015.
- [I-D.delcarpio-6lo-wlanah]
Vega, L., Robles, I., and R. Morabito, "IPv6 over 802.11ah", draft-delcarpio-6lo-wlanah-01 (work in progress), October 2015.
- [I-D.ietf-6lo-ap-nd]
Thubert, P., Sarikaya, B., and M. Sethi, "Address Protected Neighbor Discovery for Low-power and Lossy Networks", draft-ietf-6lo-ap-nd-06 (work in progress), February 2018.
- [I-D.ietf-6lo-nfc]
Choi, Y., Hong, Y., Youn, J., Kim, D., and J. Choi, "Transmission of IPv6 Packets over Near Field Communication", draft-ietf-6lo-nfc-09 (work in progress), January 2018.
- [I-D.ietf-6man-rs-refresh]
Nordmark, E., Yourtchenko, A., and S. Krishnan, "IPv6 Neighbor Discovery Optional RS/RA Refresh", draft-ietf-6man-rs-refresh-02 (work in progress), October 2016.
- [I-D.ietf-6tisch-architecture]
Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", draft-ietf-6tisch-architecture-13 (work in progress), November 2017.

- [I-D.ietf-6tisch-terminology]
Palattella, M., Thubert, P., Watteyne, T., and Q. Wang,
"Terminology in IPv6 over the TSCH mode of IEEE
802.15.4e", draft-ietf-6tisch-terminology-09 (work in
progress), June 2017.
- [I-D.ietf-bier-architecture]
Wijnands, I., Rosen, E., Dolganow, A., Przygienda, T., and
S. Aldrin, "Multicast using Bit Index Explicit
Replication", draft-ietf-bier-architecture-08 (work in
progress), September 2017.
- [I-D.ietf-ipv6-multilink-subnets]
Thaler, D. and C. Huitema, "Multi-link Subnet Support in
IPv6", draft-ietf-ipv6-multilink-subnets-00 (work in
progress), July 2002.
- [I-D.nordmark-6man-dad-approaches]
Nordmark, E., "Possible approaches to make DAD more robust
and/or efficient", draft-nordmark-6man-dad-approaches-02
(work in progress), October 2015.
- [I-D.popa-6lo-6loplc-ipv6-over-ieee19012-networks]
Popa, D. and J. Hui, "6LoPLC: Transmission of IPv6 Packets
over IEEE 1901.2 Narrowband Powerline Communication
Networks", draft-popa-6lo-6loplc-ipv6-over-
ieee19012-networks-00 (work in progress), March 2014.
- [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.
- [RFC3810] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener
Discovery Version 2 (MLDv2) for IPv6", RFC 3810,
DOI 10.17487/RFC3810, June 2004,
<<https://www.rfc-editor.org/info/rfc3810>>.

- [RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SECure Neighbor Discovery (SEND)", RFC 3971, DOI 10.17487/RFC3971, March 2005, <<https://www.rfc-editor.org/info/rfc3971>>.
- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, DOI 10.17487/RFC3972, March 2005, <<https://www.rfc-editor.org/info/rfc3972>>.
- [RFC4389] Thaler, D., Talwar, M., and C. Patel, "Neighbor Discovery Proxies (ND Proxy)", RFC 4389, DOI 10.17487/RFC4389, April 2006, <<https://www.rfc-editor.org/info/rfc4389>>.
- [RFC4903] Thaler, D., "Multi-Link Subnet Issues", RFC 4903, DOI 10.17487/RFC4903, June 2007, <<https://www.rfc-editor.org/info/rfc4903>>.
- [RFC4919] Kushalnagar, N., Montenegro, G., and C. Schumacher, "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals", RFC 4919, DOI 10.17487/RFC4919, August 2007, <<https://www.rfc-editor.org/info/rfc4919>>.
- [RFC5415] Calhoun, P., Ed., Montemurro, M., Ed., and D. Stanley, Ed., "Control And Provisioning of Wireless Access Points (CAPWAP) Protocol Specification", RFC 5415, DOI 10.17487/RFC5415, March 2009, <<https://www.rfc-editor.org/info/rfc5415>>.
- [RFC6275] Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, DOI 10.17487/RFC6275, July 2011, <<https://www.rfc-editor.org/info/rfc6275>>.
- [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>>.
- [RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", RFC 6830, DOI 10.17487/RFC6830, January 2013, <<https://www.rfc-editor.org/info/rfc6830>>.
- [RFC7048] Nordmark, E. and I. Gashinsky, "Neighbor Unreachability Detection Is Too Impatient", RFC 7048, DOI 10.17487/RFC7048, January 2014, <<https://www.rfc-editor.org/info/rfc7048>>.

- [RFC7102] Vasseur, JP., "Terms Used in Routing for Low-Power and Lossy Networks", RFC 7102, DOI 10.17487/RFC7102, January 2014, <<https://www.rfc-editor.org/info/rfc7102>>.
- [RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <<https://www.rfc-editor.org/info/rfc7217>>.
- [RFC7428] Brandt, A. and J. Buron, "Transmission of IPv6 Packets over ITU-T G.9959 Networks", RFC 7428, DOI 10.17487/RFC7428, February 2015, <<https://www.rfc-editor.org/info/rfc7428>>.
- [RFC7559] Krishnan, S., Anipko, D., and D. Thaler, "Packet-Loss Resiliency for Router Solicitations", RFC 7559, DOI 10.17487/RFC7559, May 2015, <<https://www.rfc-editor.org/info/rfc7559>>.
- [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>>.
- [RFC7772] Yourtchenko, A. and L. Colitti, "Reducing Energy Consumption of Router Advertisements", BCP 202, RFC 7772, DOI 10.17487/RFC7772, February 2016, <<https://www.rfc-editor.org/info/rfc7772>>.
- [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>>.

11.3. External Informative References

[IEEEstd8021]

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

[IEEEstd80211]

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

[IEEEstd802151]

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

[IEEEstd802154]

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

Appendix A. Requirements

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

A.1. Requirements Related to Mobility

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

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

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

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

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

A.2. Requirements Related to Routing Protocols

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

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

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

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

Related requirements are:

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

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

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

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

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

Related requirements are:

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

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

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

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

A.4. Requirements Related to Proxy Operations

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

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

Related requirements are:

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

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

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

A.5. Requirements Related to Security

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

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

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

Related requirements are:

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

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

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

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

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

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

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

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

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

A.6. Requirements Related to Scalability

Use cases from Automatic Meter Reading (AMR, collection tree operations) and Advanced Metering Infrastructure (AMI, bi-directional communication to the meters) indicate the needs for a large number of LLN nodes pertaining to a single RPL DODAG (e.g. 5000) and connected to the 6LBR over a large number of LLN hops (e.g. 15).

Related requirements are:

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

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

Author's Address

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

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

6Lo Working Group
Internet-Draft

Intended status: Standards Track
Expires: January 3, 2019

C. Gomez
S. Darroudi
Universitat Politecnica de Catalunya
T. Savolainen
DarkMatter
M. Spoerk
Graz University of Technology
July 2, 2018

IPv6 Mesh over BLUETOOTH(R) Low Energy using IPSP
draft-ietf-6lo-blemesh-03

Abstract

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

Status of This Memo

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

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

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

This Internet-Draft will expire on January 3, 2019.

Copyright Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

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

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

Table of Contents

1. Introduction	2
1.1. Terminology and Requirements Language	3
2. Bluetooth LE Networks and the IPSP	3
3. Specification of IPv6 mesh over Bluetooth LE networks	4
3.1. Protocol stack	4
3.2. Subnet model	4
3.3. Link model	5
3.3.1. Stateless address autoconfiguration	5
3.3.2. Neighbor Discovery	5
3.3.3. Header compression	7
3.3.4. Unicast and multicast mapping	8
4. IANA Considerations	8
5. Security Considerations	8
6. Contributors	8
7. Acknowledgements	9
8. Appendix	9
9. References	12
9.1. Normative References	12
9.2. Informative References	13
Authors' Addresses	13

1. Introduction

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

and would fail to enable IPv6 over mesh networks composed of Bluetooth LE links. This document specifies the mechanisms needed to enable IPv6 over mesh networks composed of Bluetooth LE links. This specification also allows to run IPv6 over Bluetooth LE star topology networks, albeit without all the topology-specific optimizations contained in RFC 7668.

1.1. Terminology and Requirements Language

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

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

2. Bluetooth LE Networks and the IPSP

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

This document assumes a mesh network composed of Bluetooth LE links, where link layer connections are established between neighboring IPv6-enabled devices (see Section 3.3.2, item 3.b)). The IPv6 forwarding devices of the mesh have to implement both Node and Router roles, while simpler leaf-only nodes can implement only the Node role. In an IPv6-enabled mesh of Bluetooth LE links, a node is a neighbor of another node, and vice versa, if a link layer connection has been established between both by using the IPSP functionality for discovery and link layer connection establishment for IPv6 packet transport.

3. Specification of IPv6 mesh over Bluetooth LE networks

3.1. Protocol stack

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

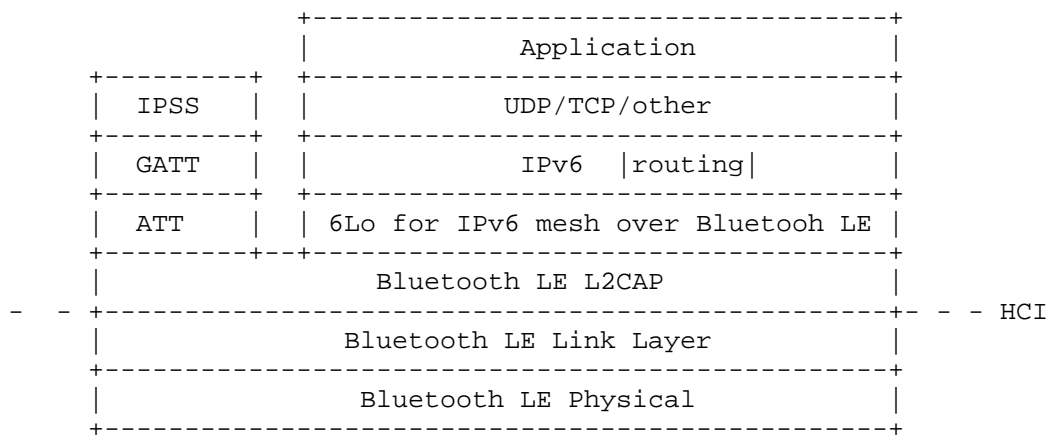


Figure 1: Protocol stack for IPv6 mesh over Bluetooth LE.

3.2. Subnet model

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




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

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

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

3.3. Link model

3.3.1. Stateless address autoconfiguration

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

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

3.3.2. Neighbor Discovery

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

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

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

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

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

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

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

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

3.3.3. Header compression

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

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

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

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

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

When a router transmits a packet to a neighboring 6LN, with a non-link-local destination address, the router MUST fully elide the destination IPv6 address if the destination address is the latest registered by the 6LN with ARO for the indicated context (DAC=1, DAM=11). If the destination address is a non-link-local address and not the latest registered, then the 6LN MUST either include the IID

part fully in-line (DAM=01) or, if the first 48-bits of the IID match to the latest registered address, then elide those 48-bits (DAM=10).

3.3.4. Unicast and multicast mapping

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

4. IANA Considerations

There are no IANA considerations related to this document.

5. Security Considerations

The security considerations in RFC 7668 apply.

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

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

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

6. Contributors

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

7. Acknowledgements

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

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

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

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

8. Appendix

This appendix provides an example of Bluetooth LE connection establishment and use of IPSP roles in an IPv6-enabled mesh of Bluetooth LE links that uses dynamic configuration. The example follows text in Section 3.3.2, item 3.b).

The example assumes a network with one 6LBR, two 6LRs and three 6LNs, as shown in Figure 3. Connectivity between the 6LNs and the 6LBR is only possible via the 6LRs.

The following text describes the different steps as time evolves, in the example. Note that other sequences of events that may lead to the same final scenario are also possible.

At the beginning, the 6LBR starts running as an IPSP Router, whereas the rest of devices are not yet initialized (Step 1). Next, the 6LRs start running as IPSP Nodes, i.e., they use Bluetooth LE advertisement packets to announce their presence and support of IPv6 capabilities (Step 2). The 6LBR (already running as an IPSP Router) discovers the presence of the 6LRs and establishes one Bluetooth LE connection with each 6LR (Step 3). After establishment of those link layer connections (and after reception of Router Advertisements from the 6LBR), Step 4, the 6LRs start operating as routers, and also initiate the IPSP Router role (note: whether the IPSP Node role is kept running simultaneously is an implementation decision). Then, 6LNs start running the IPSP Node role (Step 5). Finally, the 6LRs discover presence of the 6LNs and establish connections with the latter (Step 6).

Step 1

6LBR
(IPSP: Router)

6LR 6LR
(not initialized) (not initialized)

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


Step 2

6LBR
(IPSP: Router)


6LR 6LR
(IPSP: Node) (IPSP: Node)

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


Step 3



Step 4



Step 5



Step 6




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

9. References

9.1. Normative References

- [BTCorev4.1] Bluetooth Special Interest Group, "Bluetooth Core Specification Version 4.1", December 2013, <<https://www.bluetooth.org/en-us/specification/adopted-specifications>>.
- [IPSP] Bluetooth Special Interest Group, "Bluetooth Internet Protocol Support Profile Specification Version 1.0.0", December 2014, <<https://www.bluetooth.org/en-us/specification/adopted-specifications>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<https://www.rfc-editor.org/info/rfc6282>>.
- [RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <<https://www.rfc-editor.org/info/rfc6775>>.
- [RFC7668] Nieminen, J., Savolainen, T., Isomaki, M., Patil, B., Shelby, Z., and C. Gomez, "IPv6 over BLUETOOTH(R) Low Energy", RFC 7668, DOI 10.17487/RFC7668, October 2015, <<https://www.rfc-editor.org/info/rfc7668>>.

9.2. Informative References

- [RFC4903] Thaler, D., "Multi-Link Subnet Issues", RFC 4903, DOI 10.17487/RFC4903, June 2007, <<https://www.rfc-editor.org/info/rfc4903>>.
- [RFC7416] Tsao, T., Alexander, R., Dohler, M., Daza, V., Lozano, A., and M. Richardson, Ed., "A Security Threat Analysis for the Routing Protocol for Low-Power and Lossy Networks (RPLs)", RFC 7416, DOI 10.17487/RFC7416, January 2015, <<https://www.rfc-editor.org/info/rfc7416>>.

Authors' Addresses

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

Email: carlesgo@entel.upc.edu

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

Email: sm.darroudi@entel.upc.edu

Teemu Savolainen
DarkMatter LLC

Email: teemu.savolainen@darkmatter.ae

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

Email: michael.spoerk@tugraz.at

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

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

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

Abstract

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

Status of This Memo

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

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

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

This Internet-Draft will expire on February 24, 2021.

Copyright Notice

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

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

Table of Contents

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

1. Introduction

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

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

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

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

2. Conventions and Terminology

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

3. Overview of Near Field Communication Technology

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

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

3.1. Peer-to-peer Mode of NFC

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

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

3.2. Protocol Stack of NFC

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

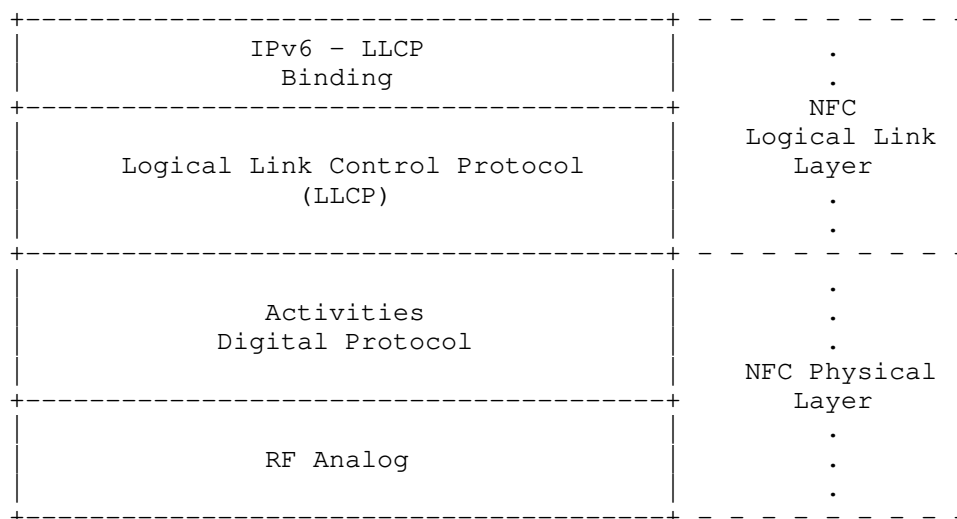


Figure 1: Protocol Stack of NFC

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

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

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

3.3. NFC-enabled Device Addressing

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

3.4. MTU of NFC Link Layer

As mentioned in Section 3.2, when an IPv6 packet is transmitted, the packet MUST be passed down to LLC of NFC and transported to an 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 may announce a larger MIU for a data link connection by transmitting an optional Maximum Information Unit Extension (MIUX) parameter within

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

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

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

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

Figure 2: Example of MIUX Parameter TLV

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

4. Specification of IPv6 over NFC

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

4.1. Protocol Stack

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




Figure 3: Protocol Stack for IPv6 over NFC

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

4.2. Stateless Address Autoconfiguration

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

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




Figure 4: IID from NFC-enabled device

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

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

4.3. IPv6 Link-Local Address

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




Figure 5: IPv6 link-local address in NFC

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

4.4. Neighbor Discovery

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

- o When an NFC-enabled 6LN is directly connected to an NFC-enabled 6LBR, the NFC 6LN MUST register its address with the 6LBR by sending a Neighbor Solicitation (NS) message with the Extended Address Registration Option (EARO) [RFC8505], and process the Neighbor Advertisement (NA) accordingly. In addition, when the 6LN and 6LBR are directly connected, DHCPv6 is used for address assignment. Therefore, Duplicate Address Detection (DAD) is not necessary between them.
- o When two or more NFC devices are connected, there are two cases. One is that three or more NFC devices are linked with multi-hop connections, and the other is that they meet within a single hop range. Two NFC devices might still talk to each other (point-to-point topology), but one of them may be connected to the Internet. In a case of multi-hop topology, devices which have two or more

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

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

4.5. Dispatch Header

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




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

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

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

Figure 7: Dispatch Values

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

4.6. Header Compression

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

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

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




Figure 8: NFC short address format

4.7. Fragmentation and Reassembly Considerations

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

4.8. Unicast and Multicast Address Mapping

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

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




Figure 9: Unicast address mapping

Option fields:

Type:

- 1: for Source Link-layer address.
- 2: for Target Link-layer address.

Length:

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

NFC address:

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

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

5. Internet Connectivity Scenarios

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

5.1. NFC-enabled Device Network Connected to the Internet

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




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

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

5.2. Isolated NFC-enabled Device Network

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

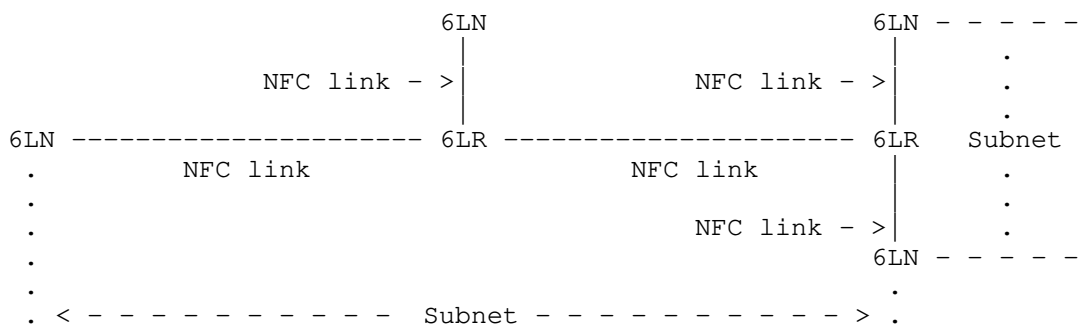


Figure 11: Isolated NFC-enabled device network

6. IANA Considerations

There are no IANA considerations related to this document.

7. Security Considerations

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

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

8. Acknowledgements

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

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

9. Normative References

[ECMA-340]

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

[LLCP-1.3]

"NFC Logical Link Control Protocol version 1.3", NFC Forum Technical Specification , March 2016.

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

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

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

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

- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<https://www.rfc-editor.org/info/rfc4944>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<https://www.rfc-editor.org/info/rfc6282>>.
- [RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <<https://www.rfc-editor.org/info/rfc6775>>.
- [RFC7136] Carpenter, B. and S. Jiang, "Significance of IPv6 Interface Identifiers", RFC 7136, DOI 10.17487/RFC7136, February 2014, <<https://www.rfc-editor.org/info/rfc7136>>.
- [RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <<https://www.rfc-editor.org/info/rfc7217>>.
- [RFC7400] Bormann, C., "6LoWPAN-GHC: Generic Header Compression for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 7400, DOI 10.17487/RFC7400, November 2014, <<https://www.rfc-editor.org/info/rfc7400>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8505] Thubert, P., Ed., Nordmark, E., Chakrabarti, S., and C. Perkins, "Registration Extensions for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Neighbor Discovery", RFC 8505, DOI 10.17487/RFC8505, November 2018, <<https://www.rfc-editor.org/info/rfc8505>>.

Authors' Addresses

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

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

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

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

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

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

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

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

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

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

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

P. Thubert, Ed.
Cisco
E. Nordmark
Zededa
S. Chakrabarti
Verizon
C. Perkins
Futurewei
June 19, 2018

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

Abstract

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

Status of This Memo

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

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

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

This Internet-Draft will expire on December 21, 2018.

Copyright Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

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

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

Table of Contents

1. Introduction	3
2. Terminology	4
2.1. BCP 14	4
2.2. References	4
2.3. Acronym Definitions	4
2.4. New Terms	5
3. Applicability of Address Registration Options	6
4. Extended Neighbor Discovery Options and Messages	7
4.1. Extended Address Registration Option (EARO)	7
4.2. Extended Duplicate Address Message Formats	11
4.3. Extensions to the Capability Indication Option	12
5. Updating RFC 6775	13
5.1. Extending the Address Registration Option	14
5.2. Transaction ID	16
5.2.1. Comparing TID values	16
5.3. Registration Ownership Verifier (ROVR)	17
5.4. Extended Duplicate Address Messages	19
5.5. Registering the Target Address	19
5.6. Link-Local Addresses and Registration	20
5.7. Maintaining the Registration States	21
6. Backward Compatibility	23
6.1. Signaling EARO Support	23
6.2. RFC6775-only 6LN	24
6.3. RFC6775-only 6LR	24
6.4. RFC6775-only 6LBR	24
7. Security Considerations	25
8. Privacy Considerations	26
9. IANA Considerations	27
9.1. ARO Flags	27
9.2. EARO I-Field	28
9.3. ICMP Codes	28
9.4. New ARO Status values	29
9.5. New 6LoWPAN Capability Bits	30
10. Acknowledgments	31
11. References	31
11.1. Normative References	31
11.2. Terminology Related References	32
11.3. Informative References	32

11.4. External Informative References	35
Appendix A. Applicability and Requirements Served (Not Normative)	36
Appendix B. Requirements (Not Normative)	37
B.1. Requirements Related to Mobility	37
B.2. Requirements Related to Routing Protocols	38
B.3. Requirements Related to the Variety of Low-Power Link types	39
B.4. Requirements Related to Proxy Operations	40
B.5. Requirements Related to Security	40
B.6. Requirements Related to Scalability	42
B.7. Requirements Related to Operations and Management	42
B.8. Matching Requirements with Specifications	43
Authors' Addresses	44

1. Introduction

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

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

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

These features satisfy requirements as listed in Appendix B.

2. Terminology

2.1. BCP 14

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

2.2. References

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

- o "Neighbor Discovery for IP version 6" [RFC4861],
- o "IPv6 Stateless Address Autoconfiguration" [RFC4862],
- o "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals" [RFC4919],
- o "Problem Statement and Requirements for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing" [RFC6606], and
- o "Neighbor Discovery Optimization for Low-power and Lossy Networks" [RFC6775],

2.3. Acronym Definitions

This document uses the following acronyms:

6BBR: 6LoWPAN Backbone Router

6LBR: 6LoWPAN Border Router

6LN: 6LoWPAN Node

6LR: 6LoWPAN Router

6CIO: Capability Indication Option

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

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

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

DAD: Duplicate Address Detection

DODAG: Destination-Oriented Directed Acyclic Graph

LLN: Low-Power and Lossy Network

NA: Neighbor Advertisement

NCE: Neighbor Cache Entry

ND: Neighbor Discovery

NDP: Neighbor Discovery Protocol

NS: Neighbor Solicitation

ROVR: Registration Ownership Verifier (pronounced rover)

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

RA: Router Advertisement

RS: Router Solicitation

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

2.4. New Terms

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

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

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

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

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

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

Registered Address: An address registered for the Registered Node.

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

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

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

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

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

3. Applicability of Address Registration Options

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

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

Examples include:

- o the router running out of space;
- o a registration bearing a stale sequence number which could happen if the host moves after the registration was placed;
- o a host misbehaving and attempting to register an invalid address such as the unspecified address [RFC4291];
- o a host using an address that is not topologically correct on that link.

In such cases the host will receive an error to help diagnose the issue and may retry, possibly with a different address, and possibly registering to a different router, depending on the returned error.

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

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

4. Extended Neighbor Discovery Options and Messages

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

4.1. Extended Address Registration Option (EARO)

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

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

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

Section 5.1 discusses those changes in depth.

The format of the EARO is shown in Figure 1:




Figure 1: EARO Option Format

Option Fields:

- Type:** 33
- Length:** 8-bit unsigned integer. The length of the option in units of 8 bytes.
- Status:** 8-bit unsigned integer. Indicates the status of a registration in the NA response. MUST be set to 0 in NS messages. See Table 1 below.
- Opaque:** An octet opaque to ND; the 6LN MAY pass it transparently to another process. It MUST be set to zero when not used.
- Rsvd (Reserved):** This field is unused. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.
- I:** Two-bit Integer: A value of zero indicates that the Opaque field carries an abstract index that is used to decide in which routing topology the address is expected to be injected. In that case, the Opaque field is passed to a routing process with the indication that it carries topology information, and the value of 0 indicates default. All other values of "I" are reserved and MUST NOT be used.

- R: The Registering Node sets the 'R' flag to request reachability services for the registered address from a Routing Registrar.
- T: One-bit flag. Set if the next octet is used as a TID.
- TID: One-byte unsigned integer; a Transaction ID that is maintained by the node and incremented with each transaction of one or more registrations performed at the same time to one or more 6LRs. This field MUST be ignored if the 'T' flag is not set.
- Registration Lifetime: 16-bit integer; expressed in minutes. A value of 0 indicates that the registration has ended and that the associated state MUST be removed.
- Registration Ownership Verifier (ROVR): Enables the correlation between multiple attempts to register a same IPv6 Address. The ROVR size MUST be 64 bits when backward compatibility is needed; otherwise the size MAY be 128, 192, or 256 bits.

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

Table 1: EARO Status

4.2. Extended Duplicate Address Message Formats

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

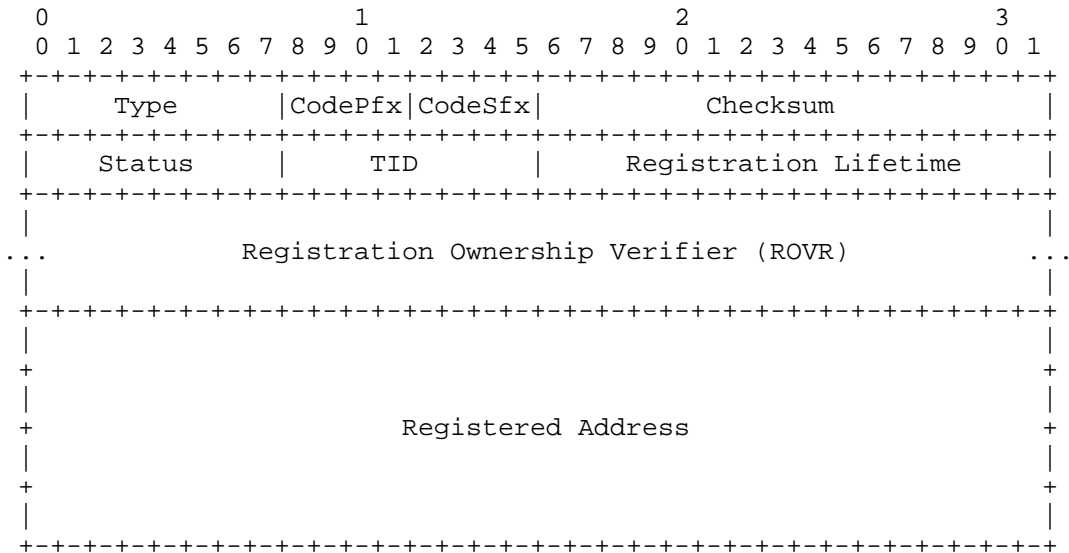


Figure 2: Duplicate Address Messages Format

Modified Message Fields:

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

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

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

4.3. Extensions to the Capability Indication Option

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

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

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

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




Figure 3: New Capability Bits in the 6CIO

Option Fields:

Type: 36

L: Node is a 6LR.

B: Node is a 6LBR.

P: Node is a Routing Registrar.

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

D: 6LBR supports EDAR and EDAC messages.

5. Updating RFC 6775

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

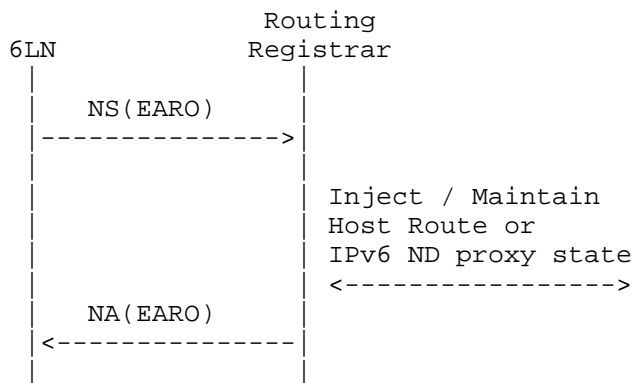


Figure 4: (Re-)Registration Flow

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

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

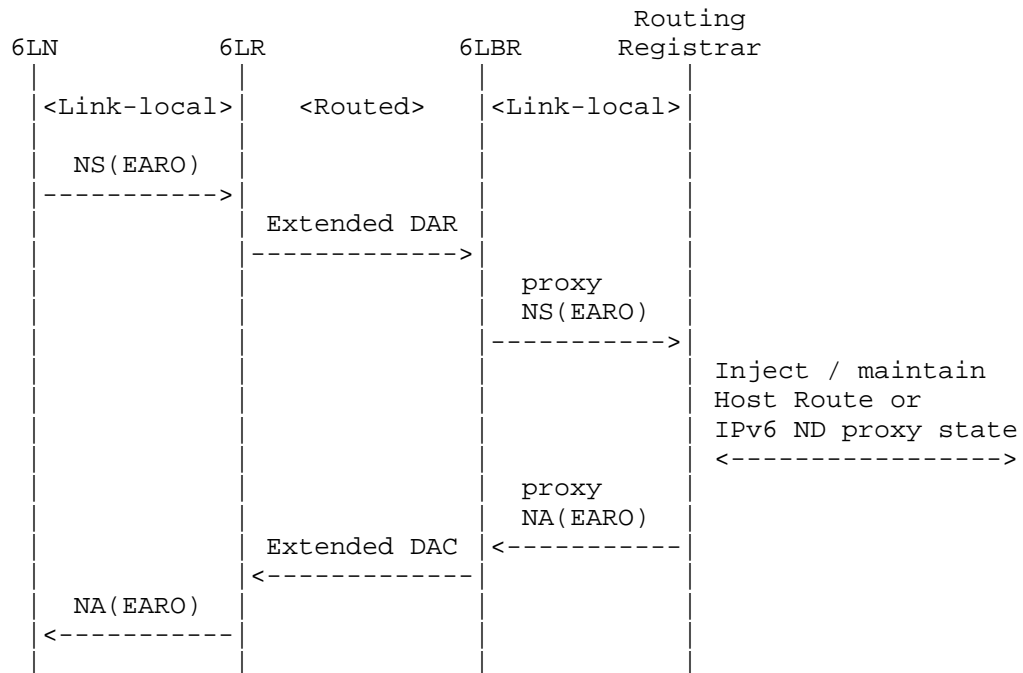


Figure 5: (Re-)Registration Flow

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

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

5.1. Extending the Address Registration Option

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

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

- o The Target Address in the NS containing the EARO is now the field that indicates the address that is being registered, as opposed to the Source Address field as specified in [RFC6775] (see Section 5.5). This change enables a 6LBR to send a proxy registration for a 6LN's address to a Routing Registrar, and also avoids in most cases the use of an address as source address before it is registered.
- o The EUI-64 field in the ARO Option is renamed Registration Ownership Verifier (ROVR) and is not required to be derived from a MAC address (see Section 5.3).
- o The option Length MAY be different than 2 and take a value between 3 and 5, in which case the EARO is not backward compatible with an ARO. The increase of size corresponds to a larger ROVR field, so the size of the ROVR is inferred from the option Length.
- o A new Opaque field is introduced to carry opaque information in case the registration is relayed to another process, e.g., to be advertised by a routing protocol. A new "I" field provides a type for the opaque information, and indicates the other process to which the 6LN passes the opaque value. A value of Zero for I indicates topological information to be passed to a routing process if the registration is redistributed. In that case, a value of Zero for the Opaque field is backward-compatible with the reserved fields that are overloaded, and the meaning is to use the default topology.
- o This document specifies a new flag in the EARO, the 'R' flag. If the 'R' flag is set, the Registering Node requests the 6LR to ensure reachability for the Registered Address, e.g., by means of routing or proxying ND. Conversely, when it is not set, the 'R' flag indicates that the Registering Node is a router, and that it will advertise reachability to the Registered Address via a routing protocol (such as RPL [RFC6550]).
- o A node that supports this specification MUST provide a Transaction ID (TID) field in the EARO, and set the 'T' flag to indicate the presence of the TID (see Section 5.2).
- o Finally, this specification introduces new status codes to help diagnose the cause of a registration failure (see Table 1).

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

5.2. Transaction ID

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

5.2.1. Comparing TID values

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

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

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

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

For a given sequence counter,

1. The sequence counter SHOULD be initialized to an implementation defined value which is 128 or greater prior to use. A recommended value is 240 ($256 - \text{SEQUENCE_WINDOW}$).
2. When a sequence counter increment would cause the sequence counter to increment beyond its maximum value, the sequence counter MUST wrap back to zero. When incrementing a sequence counter greater than or equal to 128, the maximum value is 255. When incrementing a sequence counter less than 128, the maximum value is 127.

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

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

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

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

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

1. If the absolute magnitude of difference between the two sequence counters is less than or equal to `SEQUENCE_WINDOW`, then a comparison as described in [RFC1982] is used to determine the relationships greater than, less than, and equal.
2. If the absolute magnitude of difference of the two sequence counters is greater than `SEQUENCE_WINDOW`, then a desynchronization has occurred and the two sequence numbers are not comparable.

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

5.3. Registration Ownership Verifier (ROVR)

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

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

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

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

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

5.4. Extended Duplicate Address Messages

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

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

5.5. Registering the Target Address

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

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

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

In order to enable the latter operation, this specification changes the behavior of the 6LN and the 6LR so that the Registered Address is found in the Target Address field of the NS and NA messages as opposed to the Source Address field. With this convention, a TLLA option indicates the link-layer address of the 6LN that owns the address.

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

5.6. Link-Local Addresses and Registration

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

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

An exchange between two nodes using Link-Local Addresses implies that they are reachable over one hop. A node MUST register a Link-Local Address to a 6LR in order to obtain further reachability by way of that 6LR, and in particular to use the Link-Local Address as source address to register other addresses, e.g., global addresses.

If there is no collision with a previously registered address, then the Link-Local Address is unique from the standpoint of this 6LR and the registration is not a duplicate. Two different 6LRs might claim the same Link-Local Address but different link-layer addresses. In that case, a 6LN MUST only interact with at most one of the 6LRs.

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

When sending an NS(EARO) to a 6LR, a 6LN MUST use a Link-Local Address as the source address of the registration, whatever the type of IPv6 address that is being registered. That Link-Local Address MUST be either an address that is already registered to the 6LR, or the address that is being registered.

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

When a Registering Node does not have an already-registered Address, it MUST register a Link-Local Address, using it as both the Source and the Target Address of an NS(EARO) message. In that case, it is

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

A 6LR that supports this specification replies with an NA(EARO), setting the appropriate status. Since there is no exchange of EDAR or EDAC messages for Link-Local Addresses, the 6LR may answer immediately to the registration of a Link-Local Address, based solely on its existing state and the Source Link-Layer Option that is placed in the NS(EARO) message as required in [RFC6775].

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

5.7. Maintaining the Registration States

This section discusses protocol actions that involve the Registering Node, the 6LR, and the 6LBR. It must be noted that the portion that deals with a 6LBR only applies to those addresses that are registered to it; as discussed in Section 5.6, this is not the case for Link-Local Addresses. The registration state includes all data that is stored in the router relative to that registration, in particular, but not limited to, an NCE. 6LBRs and Routing Registrars may store additional registration information and use synchronization protocols that are out of scope of this document.

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

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

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

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

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

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

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

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

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

6. Backward Compatibility

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

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

6.1. Signaling EARO Support

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

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

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

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

6.2. RFC6775-only 6LN

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

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

6.3. RFC6775-only 6LR

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

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

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

6.4. RFC6775-only 6LBR

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

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

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

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

7. Security Considerations

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

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

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

The registration mechanism may be used by a rogue node to attack the 6LR or the 6LBR with a Denial-of-Service attack against the registry. It may also happen that the registry of a 6LR or a 6LBR is saturated and cannot take any more registrations, which effectively denies the requesting node the capability to use a new address. In order to alleviate those concerns, Section 5.7 provides a number of recommendations that ensure that a stale registration is removed as soon as possible from the 6LR and 6LBR. In particular, this specification recommends that:

- o A node that ceases to use an address SHOULD attempt to de-register that address from all the 6LRs to which it is registered. See Section 5.2 for the mechanism to avoid replay attacks and avoiding the use of stale registration information.
- o The Registration lifetimes SHOULD be individually configurable for each address or group of addresses. The nodes SHOULD be configured with a Registration Lifetime that reflects their expectation of how long they will use the address with the 6LR to which it is registered. In particular, use cases that involve mobility or rapid address changes SHOULD use lifetimes that are

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

- o The router (6LR or 6LBR) SHOULD be configurable so as to limit the number of addresses that can be registered by a single node, but as a protective measure only. In any case, a router MUST be able to keep a minimum number of addresses per node. That minimum depends on the type of device and ranges between 3 for a very constrained LLN and 10 for a larger device. A node may be identified by its MAC address, as long as it is not obfuscated by privacy measures. A stronger identification (e.g., by security credentials) is RECOMMENDED. When the maximum is reached, the router SHOULD use a Least-Recently-Used (LRU) algorithm to clean up the addresses, keeping at least one Link-Local Address. The router SHOULD attempt to keep one or more stable addresses if stability can be determined, e.g., because they are used over a much longer time span than other (privacy, shorter-lived) addresses.
- o In order to avoid denial of registration for the lack of resources, administrators should take great care to deploy adequate numbers of 6LRs to cover the needs of the nodes in their range, so as to avoid a situation of starving nodes. It is expected that the 6LBR that serves an LLN is a more capable node than the average 6LR, but in a network condition where it may become saturated, a particular LLN should distribute the 6LBR functionality, for instance by leveraging a high speed Backbone Link and Routing Registrars to aggregate multiple LLNs into a larger subnet.

The LLN nodes depend on a 6LBR and may use the services of a routing Registrar for their operation. A trust model MUST be put in place to ensure that only authorized devices are acting in these roles so as to avoid threats such as black-holing or bombing attack whereby an impersonated 6LBR would destroy state in the network by using the "Removed" Status code. This trust model could be at a minimum based on a Layer-2 access control, or could provide role validation as well (see Req5.1 in Appendix B.5).

8. Privacy Considerations

As indicated in Section 3, this protocol does not limit the number of IPv6 addresses that each device can form. However, to mitigate denial-of-service attacks, it can be useful as a protective measure to have a limit that is high enough not to interfere with the normal behavior of devices in the network. A host should be able to form and register any address that is topologically correct in the subnet(s) advertised by the 6LR/6LBR.

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

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

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

9. IANA Considerations

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

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

9.1. ARO Flags

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

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

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

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

Table 2: New ARO Flags

9.2. EARO I-Field

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

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

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

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

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

9.3. ICMP Codes

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

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

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

The new subregistries are to be initialized as follows:

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

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

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

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

9.4. New ARO Status values

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

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

Table 6: New ARO Status values

9.5. New 6LoWPAN Capability Bits

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

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

Table 7: New 6LoWPAN Capability Bits

10. Acknowledgments

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

11. References

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <<https://www.rfc-editor.org/info/rfc4443>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <<https://www.rfc-editor.org/info/rfc4862>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<https://www.rfc-editor.org/info/rfc6282>>.

- [RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <<https://www.rfc-editor.org/info/rfc6775>>.
- [RFC7400] Bormann, C., "6LoWPAN-GHC: Generic Header Compression for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 7400, DOI 10.17487/RFC7400, November 2014, <<https://www.rfc-editor.org/info/rfc7400>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

11.2. Terminology Related References

- [RFC4919] Kushalnagar, N., Montenegro, G., and C. Schumacher, "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals", RFC 4919, DOI 10.17487/RFC4919, August 2007, <<https://www.rfc-editor.org/info/rfc4919>>.
- [RFC6606] Kim, E., Kaspar, D., Gomez, C., and C. Bormann, "Problem Statement and Requirements for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing", RFC 6606, DOI 10.17487/RFC6606, May 2012, <<https://www.rfc-editor.org/info/rfc6606>>.

11.3. Informative References

- [I-D.chakrabarti-nordmark-6man-efficient-nd]
Chakrabarti, S., Nordmark, E., Thubert, P., and M. Wasserman, "IPv6 Neighbor Discovery Optimizations for Wired and Wireless Networks", draft-chakrabarti-nordmark-6man-efficient-nd-07 (work in progress), February 2015.
- [I-D.delcarpio-6lo-wlanah]
Vega, L., Robles, I., and R. Morabito, "IPv6 over 802.11ah", draft-delcarpio-6lo-wlanah-01 (work in progress), October 2015.

- [I-D.hou-6lo-plc]
Hou, J., Hong, Y., and X. Tang, "Transmission of IPv6 Packets over PLC Networks", draft-hou-6lo-plc-03 (work in progress), December 2017.
- [I-D.ietf-6lo-ap-nd]
Thubert, P., Sarikaya, B., and M. Sethi, "Address Protected Neighbor Discovery for Low-power and Lossy Networks", draft-ietf-6lo-ap-nd-06 (work in progress), February 2018.
- [I-D.ietf-6lo-backbone-router]
Thubert, P., "IPv6 Backbone Router", draft-ietf-6lo-backbone-router-06 (work in progress), February 2018.
- [I-D.ietf-6lo-nfc]
Choi, Y., Hong, Y., Youn, J., Kim, D., and J. Choi, "Transmission of IPv6 Packets over Near Field Communication", draft-ietf-6lo-nfc-09 (work in progress), January 2018.
- [I-D.ietf-6tisch-architecture]
Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", draft-ietf-6tisch-architecture-14 (work in progress), April 2018.
- [I-D.ietf-mboned-ieee802-mcast-problems]
Perkins, C., McBride, M., Stanley, D., Kumari, W., and J. Zuniga, "Multicast Considerations over IEEE 802 Wireless Media", draft-ietf-mboned-ieee802-mcast-problems-01 (work in progress), February 2018.
- [I-D.ietf-roll-efficient-npdao]
Jadhav, R., Thubert, P., Sahoo, R., and Z. Cao, "Efficient Route Invalidation", draft-ietf-roll-efficient-npdao-03 (work in progress), March 2018.
- [I-D.struik-lwip-curve-representations]
Struik, R., "Alternative Elliptic Curve Representations", draft-struik-lwip-curve-representations-00 (work in progress), October 2017.
- [I-D.thubert-roll-unaware-leaves]
Thubert, P., "Routing for RPL Leaves", draft-thubert-roll-unaware-leaves-05 (work in progress), May 2018.

- [RFC1958] Carpenter, B., Ed., "Architectural Principles of the Internet", RFC 1958, DOI 10.17487/RFC1958, June 1996, <<https://www.rfc-editor.org/info/rfc1958>>.
- [RFC1982] Elz, R. and R. Bush, "Serial Number Arithmetic", RFC 1982, DOI 10.17487/RFC1982, August 1996, <<https://www.rfc-editor.org/info/rfc1982>>.
- [RFC3610] Whiting, D., Housley, R., and N. Ferguson, "Counter with CBC-MAC (CCM)", RFC 3610, DOI 10.17487/RFC3610, September 2003, <<https://www.rfc-editor.org/info/rfc3610>>.
- [RFC3810] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, DOI 10.17487/RFC3810, June 2004, <<https://www.rfc-editor.org/info/rfc3810>>.
- [RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, DOI 10.17487/RFC3971, March 2005, <<https://www.rfc-editor.org/info/rfc3971>>.
- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, DOI 10.17487/RFC3972, March 2005, <<https://www.rfc-editor.org/info/rfc3972>>.
- [RFC4429] Moore, N., "Optimistic Duplicate Address Detection (DAD) for IPv6", RFC 4429, DOI 10.17487/RFC4429, April 2006, <<https://www.rfc-editor.org/info/rfc4429>>.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, DOI 10.17487/RFC4941, September 2007, <<https://www.rfc-editor.org/info/rfc4941>>.
- [RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <<https://www.rfc-editor.org/info/rfc6550>>.
- [RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <<https://www.rfc-editor.org/info/rfc7217>>.

- [RFC7428] Brandt, A. and J. Buron, "Transmission of IPv6 Packets over ITU-T G.9959 Networks", RFC 7428, DOI 10.17487/RFC7428, February 2015, <<https://www.rfc-editor.org/info/rfc7428>>.
- [RFC7668] Nieminen, J., Savolainen, T., Isomaki, M., Patil, B., Shelby, Z., and C. Gomez, "IPv6 over BLUETOOTH(R) Low Energy", RFC 7668, DOI 10.17487/RFC7668, October 2015, <<https://www.rfc-editor.org/info/rfc7668>>.
- [RFC7934] Colitti, L., Cerf, V., Cheshire, S., and D. Schinazi, "Host Address Availability Recommendations", BCP 204, RFC 7934, DOI 10.17487/RFC7934, July 2016, <<https://www.rfc-editor.org/info/rfc7934>>.
- [RFC8064] Gont, F., Cooper, A., Thaler, D., and W. Liu, "Recommendation on Stable IPv6 Interface Identifiers", RFC 8064, DOI 10.17487/RFC8064, February 2017, <<https://www.rfc-editor.org/info/rfc8064>>.
- [RFC8065] Thaler, D., "Privacy Considerations for IPv6 Adaptation-Layer Mechanisms", RFC 8065, DOI 10.17487/RFC8065, February 2017, <<https://www.rfc-editor.org/info/rfc8065>>.
- [RFC8105] Mariager, P., Petersen, J., Ed., Shelby, Z., Van de Logt, M., and D. Barthel, "Transmission of IPv6 Packets over Digital Enhanced Cordless Telecommunications (DECT) Ultra Low Energy (ULE)", RFC 8105, DOI 10.17487/RFC8105, May 2017, <<https://www.rfc-editor.org/info/rfc8105>>.
- [RFC8163] Lynn, K., Ed., Martocci, J., Neilson, C., and S. Donaldson, "Transmission of IPv6 over Master-Slave/Token-Passing (MS/TP) Networks", RFC 8163, DOI 10.17487/RFC8163, May 2017, <<https://www.rfc-editor.org/info/rfc8163>>.
- [RFC8279] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast Using Bit Index Explicit Replication (BIER)", RFC 8279, DOI 10.17487/RFC8279, November 2017, <<https://www.rfc-editor.org/info/rfc8279>>.

11.4. External Informative References

- [IEEEstd802154] IEEE, "IEEE Standard for Low-Rate Wireless Networks", IEEE Standard 802.15.4, DOI 10.1109/IEEE P802.15.4-REVD/01, June 2017, <<http://ieeexplore.ieee.org/document/7460875/>>.

[Perlman83]

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

Appendix A. Applicability and Requirements Served (Not Normative)

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

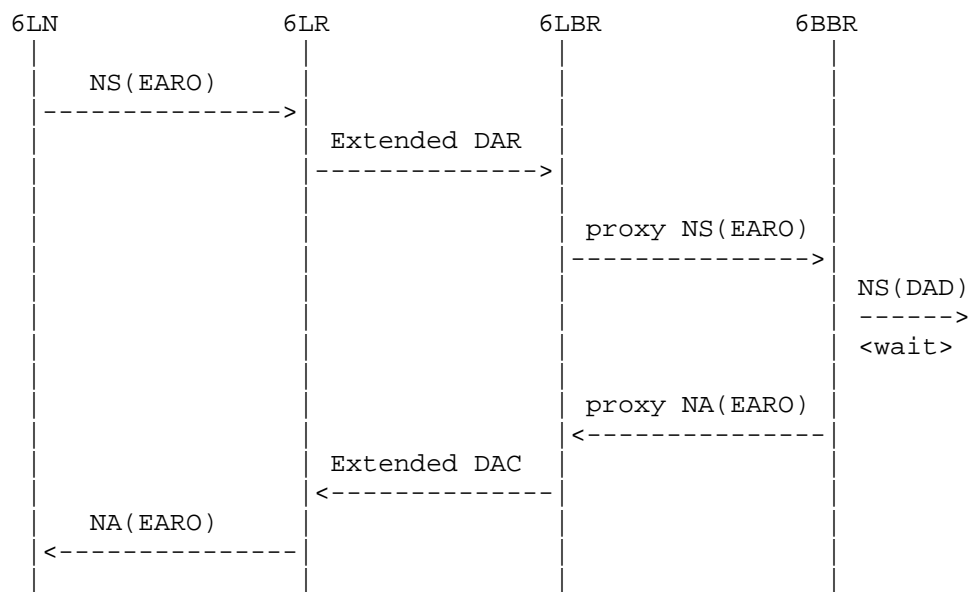


Figure 6: (Re-)Registration Flow

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

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

The term LLN is used loosely in this document, and intended to cover multiple types of WLANs and WPANs, including Low-Power IEEE Std. 802.11 networking, Bluetooth Low Energy, IEEE Std. 802.11ah, and IEEE Std. 802.15.4 wireless meshes, so as to address the requirements discussed in Appendix B.3.

This specification can be used by any wireless node to register its IPv6 addresses with a Routing Registrar and to obtain routing services including proxy-ND operations over a Backbone Link. This satisfies the the requirements expressed in Appendix B.4.

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

"Efficiency aware IPv6 Neighbor Discovery Optimizations" [I-D.chakrabarti-nordmark-6man-efficient-nd] suggests that 6LoWPAN ND [RFC6775] can be extended to other types of links beyond IEEE Std. 802.15.4 for which it was defined. The registration technique is beneficial when the Link-Layer technique used to carry IPv6 multicast packets is not sufficiently efficient in terms of delivery ratio or energy consumption in the end devices, in particular to enable energy-constrained sleeping nodes. The value of such extension is especially apparent in the case of mobile wireless nodes, to reduce the multicast operations that are related to IPv6 ND ([RFC4861], [RFC4862]) and affect the operation of the wireless medium [I-D.ietf-mboned-ieee802-mcast-problems]. This serves the scalability requirements listed in Appendix B.6.

Appendix B. Requirements (Not Normative)

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

B.1. Requirements Related to Mobility

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

and restore reachability in a timely fashion, e.g., by using some signaling upon the detection of the movement, or using a keep-alive mechanism with a period that is consistent with the application needs.

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

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

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

Req1.4: A 6LoWPAN Node SHOULD also be able to register its Address concurrently to multiple 6LRs.

B.2. Requirements Related to Routing Protocols

The point of attachment of a 6LN may be a 6LR in an LLN mesh. IPv6 routing in an LLN can be based on RPL, which is the routing protocol that was defined by the IETF for this particular purpose. Other routing protocols are also considered by Standards Development Organizations (SDO) on the basis of the expected network characteristics. It is required that a 6LN attached via ND to a 6LR indicates whether it participates in the selected routing protocol to obtain reachability via the 6LR, or whether it expects the 6LR to manage its reachability.

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

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

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

The Bit Index Explicit Replication (BIER) Architecture [RFC8279] proposes an optimized technique to enable multicast in an LLN with a very limited requirement for routing state in the nodes.

Related requirements are:

Req2.1: The ND registration method SHOULD be extended so that the 6LR is instructed whether to advertise the Address of a 6LN over the selected routing protocol and obtain reachability to that Address using the selected routing protocol.

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

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

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

6LoWPAN ND [RFC6775] was defined with a focus on IEEE Std.802.15.4 and in particular the capability to derive a unique identifier from a globally unique EUI-64 address. At this point, the 6Lo Working Group is extending the 6LoWPAN Header Compression (HC) [RFC6282] technique to other link types including ITU-T G.9959 [RFC7428], Master-Slave/Token-Passing [RFC8163], DECT Ultra Low Energy [RFC8105], Near Field Communication [I-D.ietf-6lo-nfc], IEEE Std. 802.11ah [I-D.delcarpio-6lo-wlanah], as well as Bluetooth(R) Low Energy [RFC7668], and Power Line Communication (PLC) [I-D.hou-6lo-plc] Networks.

Related requirements are:

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

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

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

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

B.4. Requirements Related to Proxy Operations

Duty-cycled devices may not be awake to answer a lookup from a node that uses IPv6 ND and may need a proxy. Additionally, the duty-cycled device may rely on the 6LBR to perform registration to the Routing Registrar.

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

Related requirements are:

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

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

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

B.5. Requirements Related to Security

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

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

In an LLN it makes sense to base security on Layer-2 security. During bootstrap of the LLN, nodes join the network after authorization by a Joining Assistant (JA) or a Commissioning Tool (CT). After joining, nodes communicate with each other via secured

links. The keys for the Layer-2 security are distributed by the JA/CT. The JA/CT can be part of the LLN or be outside the LLN. In both cases it is needed that packets are routed between JA/CT and the joining node.

Related requirements are:

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

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

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

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

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

Req5.6: The 6LoWPAN ND security mechanisms SHOULD enable the variation of CCM [RFC3610] called CCM* for use at both Layer 2 and Layer 3, and SHOULD enable the reuse of security code that has to be present on the device for upper layer security such as TLS. Algorithm agility and support for large keys (e.g., 256-bit key sizes) is also desirable, following at Layer-3 the introduction of those capabilities at Layer-2.

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

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

Req5.9: 6LoWPAN ND security mechanisms SHOULD provide a mechanism for the 6LR and the 6LBR to validate whether a new registration for a given address corresponds to the same 6LN that registered it

initially, and, if not, determine the rightful owner and deny or clean up the registration that is duplicate.

B.6. Requirements Related to Scalability

Use cases from Automatic Meter Reading (AMR, collection tree operations) and Advanced Metering Infrastructure (AMI, bi-directional communication to the meters) indicate the needs for a large number of LLN nodes pertaining to a single RPL DODAG (e.g., 5000) and connected to the 6LBR over a large number of LLN hops (e.g., 15).

Related requirements are:

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

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

B.7. Requirements Related to Operations and Management

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

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

Related requirements are:

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

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

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

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

B.8. Matching Requirements with Specifications

I-drafts/RFCs addressing requirements

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

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

Table 8: Work Addressing requirements

Authors' Addresses

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

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

Erik Nordmark
Zededa
Santa Clara, CA
United States of America

Email: nordmark@sonic.net

Samita Chakrabarti
Verizon
San Jose, CA
United States of America

Email: samitac.ietf@gmail.com

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

Email: charliep@computer.org

6Lo Working Group
Internet-Draft
Intended status: Informational
Expires: January 1, 2019

Y-G. Hong
ETRI
C. Gomez
UPC
Y-H. Choi
ETRI
AR. Santi
Huaiyin Institute of Technology
T. Aanstoot
Modio AB
S. Chakrabarti
June 30, 2018

IPv6 over Constrained Node Networks (6Lo) Applicability & Use cases
draft-ietf-6lo-use-cases-05

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 January 1, 2019.

Copyright Notice

Copyright (c) 2018 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 6LoWPAN (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 6LoWPAN [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 6LoWPAN 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 6LoWPAN 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.

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-transmitt the frame on a different frequency. The resulting "channel hopping" efficiently combats external interference and multi-path fading.

The main benefits of IEEE 802.15.4 TSCH are:

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

IEEE 802.15.4e can be used for industrial automation.

3.8. Comparison between 6lo Link layer technologies

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

	Z-Wave	BLE	DECT-ULE	MS/TP	NFC	PLC	TSCH
Usage	Home Auto-mation	Interact w/ Smart Phone	Meter Reading	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-hou-6lo-plc	RFC7554

Table 2: Comparison between 6lo Link layer technologies

4. 6lo Deployment Scenarios

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

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

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

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

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

jupiterMesh specification is based on the following technologies:

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

- o Address assignment and network configuration are specified using DHCPv6 [RFC3315]. Neighbor Discovery (ND) [RFC6775] and stateless address auto-configuration (SLAAC) are not supported.
- o The network layer consists of IPv6, 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 [I-D.ietf-lwig-energy-efficient]. Readers are expected to be familiar with [RFC7228] terminology.
- o Update firmware requirements: Most 6lo use cases will need a mechanism for updating firmware. In these cases support for over the air updates are required, probably in a broadcast mode when bandwidth is low and the number of identical devices is high.
- 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>>.

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-09 (work in progress), January 2018.

- [I-D.ietf-lwig-energy-efficient]
Gomez, C., Kovatsch, M., Tian, H., and Z. Cao, "Energy-Efficient Features of Internet of Things Protocols", draft-ietf-lwig-energy-efficient-08 (work in progress), October 2017.
- [I-D.ietf-roll-aodv-rpl]
Anamalamudi, S., Zhang, M., Sangi, A., Perkins, C., Anand, S., and B. Liu, "Asymmetric AODV-P2P-RPL in Low-Power and Lossy Networks (LLNs)", draft-ietf-roll-aodv-rpl-03 (work in progress), March 2018.
- [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., and T. Savolainen, "IPv6 Mesh over BLUETOOTH(R) Low Energy using IPSP", draft-ietf-6lo-blemesh-02 (work in progress), September 2017.
- [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.
- [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>.
- [G3-PLC]
"G3-PLC Alliance", <<http://www.g3-plc.com/home/>>.
- [NETRICITY]
"Netricity program in HomePlug Powerline Alliance",
<<http://groups.homeplug.org/tech/Netricity>>.

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

Appendix A. Other 6lo Use Case Examples

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

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

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

Variety of home devices (e.g. light dimmers/switches, plugs, thermostats, blinds/curtains and remote controls) are augmented with

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

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

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

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

Example: Use of DECT-ULE for Smart Metering

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

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

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

Example: Use of 6LoBAC in Building Automation Networks

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

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

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

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

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

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

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

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

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

A.5. Use case of PLC: Smart Grid

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

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

Example: Use of PLC for Advanced Metering Infrastructure

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

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

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

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

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

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

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

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

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

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

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

Authors' Addresses

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

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

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

Email: carlesgo@entel.upc.edu

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

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

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

Email: sangi_bahrian@yahoo.com

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

Email: take@modio.se

Samita Chakrabarti
San Jose, CA
USA

Email: samitac.ietf@gmail.com

6lo
Internet-Draft
Updates: 4944 (if approved)
Intended status: Standards Track
Expires: July 20, 2018

P. Thubert, Ed.
Cisco Systems
J. Hui
Nest Labs
January 16, 2018

LLN Fragment Forwarding and Recovery
draft-thubert-6lo-forwarding-fragments-08

Abstract

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

Status of This Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <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 20, 2018.

Copyright Notice

Copyright (c) 2018 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. Updating RFC 4944	3
3. Terminology and Referenced Work	4
4. New Dispatch types and headers	5
4.1. Recoverable Fragment Dispatch type and Header	5
4.2. RFRAG Acknowledgment Dispatch type and Header	7
5. Fragments Recovery	9
6. Forwarding Fragments	10
6.1. Upon the first fragment	11
6.2. Upon the next fragments	12
6.3. Upon the RFRAG Acknowledgments	12
7. Security Considerations	13
8. IANA Considerations	13
9. Acknowledgments	13
10. References	13
10.1. Normative References	13
10.2. Informative References	14
Appendix A. Rationale	15
Appendix B. Requirements	17
Appendix C. Considerations On Flow Control	18
Authors' Addresses	19

1. Introduction

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

"Transmission of IPv6 Packets over IEEE 802.15.4 Networks" [RFC4944] defines the original 6LoWPAN datagram fragmentation mechanism for LLNs. One critical issue with this original design is that routing an IPv6 [RFC8200] packet across a route-over mesh requires to reassemble the full packet at each hop, which may cause latency along a path and an overall buffer bloat in the network. Those undesirable effects can be alleviated by a hop-by-hop fragment forwarding technique such as the one proposed in this specification, and arguably this could be achieved without the need to define a new protocol. However, adding that capability alone to the local implementation of the original 6LoWPAN fragmentation would not address the bulk of the issues raised against it, and may create new issues like uncontrolled state in the network.

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

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

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

2. Updating RFC 4944

This specification 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. New dispatch types are defined in Section 4.

3. Terminology and Referenced Work

Past experience with fragmentation has shown that miss-associated or lost fragments can lead to poor network behavior and, occasionally, trouble at application layer. The reader is encouraged to read "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.

This specification uses the following terms:

6LoWPAN endpoints

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

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

4. New Dispatch types and headers

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

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

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

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

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

Figure 1: Additional Dispatch Value Bit Patterns

4.1. Recoverable Fragment Dispatch type and Header

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

The first fragment is recognized by a sequence of 0; it carries its fragment_size and the datagram_size of the compressed packet, whereas

the other fragments carry their `fragment_size` and `fragment_offset`. The last fragment for a datagram is recognized when its `fragment_offset` and its `fragment_size` add up to the `datagram_size`.

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




Figure 2: RFRAG Dispatch type and Header

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

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

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

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

Fragment_offset: 11 bit unsigned integer;

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

- + When the `Sequence` is not 0, this field indicates the offset of the fragment in the compressed form. The fragment should be forwarded based on an existing label Switched Path (LSP) as described in Section 6.2, or silently dropped if none is found.

- + When the Sequence is 0, denoting the first fragment of a datagram, this field is overloaded to indicate the `total_size` of the compressed packet, to help the receiver allocate an adapted buffer for the reception and reassembly operations. This format limits the maximum MTU on a 6LoWPAN link to 2047 bytes, but 1280 bytes is the recommended value to avoid issues with IPV6 Path MTU Discovery [RFC8201]. The fragment should be routed based on the destination IPv6 address, and an LSP state should be installed as described in Section 6.1.
- * When set to 0, this field indicates an abort condition and all state regarding the datagram should be cleaned up once the processing of the fragment is complete; the processing of the fragment depends on whether there is an LSP already established for this datagram, and the next hop is still reachable:
 - + if an LSP already exists and is not broken, the fragment is to be forwarded along that LSP as described in Section 6.2, 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 but no state is conserved afterwards.

If the fragment cannot be forwarded or routed, then it is silently dropped.

4.2. 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: Expanding 3 octets encoding

Figure 5: RFRAG Acknowledgment Dispatch type and Header

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.

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.

5. Fragments Recovery

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

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

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

The receiver MAY issue unsolicited acknowledgments. An unsolicited acknowledgment signals to the sender endpoint that it can resume sending if it had reached its maximum number of outstanding

fragments. Another use is to inform that the reassembling endpoint has cancelled the process of an individual datagram. Note that acknowledgments might consume precious resources so the use of unsolicited acknowledgments should be configurable and not enabled by default.

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

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

When the sender decides that a packet should be dropped and the fragmentation process canceled, it sends a pseudo fragment with the `fragment_offset`, `sequence` and `fragment_size` all set to 0, and no data. Upon reception of this message, the receiver should clean up all resources for the packet associated to the `datagram_tag`. If an acknowledgment is requested, the receiver responds with a NULL bitmap.

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

6. Forwarding Fragments

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

This specification enables intermediate routers to forward fragments with no intermediate reconstruction of the entire packet. The first fragment carries the IP header and it is routed all the way from the fragmenting end point to the reassembling end point. Upon the first fragment, the routers along the path install a label-switched path

(LSP), and the following fragments are label-switched along that path. As a consequence, alternate routes not possible for individual fragments. The datagram_tag is used to carry the label, that is swapped at each hop. All fragments follow the same path and fragments are delivered in the order at which they are sent.

6.1. 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. Say the first fragment has:

- o Source IPv6 address = IP_A (maybe hops away)
- o Destination IPv6 address = IP_B (maybe hops away)
- o Source MAC = MAC_previous
- o Datagram_tag= DT_previous

The intermediate router that forwards individual fragments performs the following action:

1. a route lookup to get the Next hop IPv6 towards IP_B, which resolves as IP_next.
2. a MAC address resolution to get the MAC address associated to IP_next, which resolves as MAC_next

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

1. cleans up any leftover resource associated to the tuple (MAC_previous, DT_previous)
2. allocates a new label for that flow, DT_next, from a Least Recently Used pool or some similar procedure.
3. allocates an abstract label-swap entry indexed by (MAC_previous, DT_previous) that contains (MAC_next, DT_next)
4. allocates a reflective abstract label-swap structure indexed by (MAC_next, DT_next) that contains (MAC_previous, DT_previous); this enables the reverse MPLS switching operation that is used to route the RFRAG-ACK.
5. change the source MAC address from MAC_prev to MAC_self

6. change the destination MAC address to from MAC_self to MAC_next
7. Swaps the datagram_tag to DT_next

At this point the router is all set and can forward the fragment to next.

6.2. Upon the next fragments

Upon next fragments (that are not first fragment), the router expects to have already installed a label-swap structure indexed by (MAC_previous, DT_previous). The router:

1. looks up the label-swap entry for (MAC_previous, DT_previous), which resolves as (MAC_next, DT_next)
2. swaps the MAC info to from self to MAC_next;
3. Swaps the datagram_tag to DT_next

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

- o MAC info set to from self to MAC_previous as found in the fragment
- o The datagram_tag set to DT_previous
- o Null bitmap to indicate the error

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

6.3. Upon the RFRAG Acknowledgments

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

1. looks up the label-swap entry for (MAC_next, DT_next), which resolves as (MAC_previous, DT_previous)
2. swaps the MAC info to from self to MAC_previous;
3. Swaps the datagram_tag to DT_previous

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

If the label-swap entry for (MAC_next, DT_next) is not found, it MUST silently drop the packet.

If the RFRAG-ACK indicates either an error (NULL bitmap) or that the fragment was entirely received (FULL bitmap), the router schedules the label-swap entries for recycling. If the RFRAG-ACK is lost on the way back, the source may retry the last fragment, which will result as an error RFRAG-ACK from the first router on the way that has already cleaned up.

7. Security Considerations

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

8. IANA Considerations

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

9. Acknowledgments

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

10. References

10.1. Normative References

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

10.2. Informative References

- [I-D.ietf-6tisch-architecture] Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", draft-ietf-6tisch-architecture-13 (work in progress), November 2017.
- [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>>.
- [RFC7554] Watteyne, T., Ed., Palattella, M., and L. Grieco, "Using IEEE 802.15.4e Time-Slotted Channel Hopping (TSCH) in the Internet of Things (IoT): Problem Statement", RFC 7554, DOI 10.17487/RFC7554, May 2015, <<https://www.rfc-editor.org/info/rfc7554>>.
- [RFC7567] Baker, F., Ed. and G. Fairhurst, Ed., "IETF Recommendations Regarding Active Queue Management", BCP 197, RFC 7567, DOI 10.17487/RFC7567, July 2015, <<https://www.rfc-editor.org/info/rfc7567>>.
- [RFC8085] Eggert, L., Fairhurst, G., and G. Shepherd, "UDP Usage Guidelines", BCP 145, RFC 8085, DOI 10.17487/RFC8085, March 2017, <<https://www.rfc-editor.org/info/rfc8085>>.
- [RFC8087] Fairhurst, G. and M. Welzl, "The Benefits of Using Explicit Congestion Notification (ECN)", RFC 8087, DOI 10.17487/RFC8087, March 2017, <<https://www.rfc-editor.org/info/rfc8087>>.
- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.
- [RFC8201] McCann, J., Deering, S., Mogul, J., and R. Hinden, Ed., "Path MTU Discovery for IP version 6", STD 87, RFC 8201, DOI 10.17487/RFC8201, July 2017, <<https://www.rfc-editor.org/info/rfc8201>>.

Appendix A. Rationale

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

The list of those usages includes:

Towards the LLN node:

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

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

From the LLN node:

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

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

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

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

When a fragment is lost in transmission, the lack of recovery in the original fragmentation system of RFC 4944 implies that all fragments are resent, further contributing to the congestion that caused the initial loss, and potentially leading to congestion collapse.

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

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

IPv4 fragments. At the same time, the use of radios increases the probability of transmission loss and Mesh-Under techniques compound that risk over multiple hops.

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

Appendix B. Requirements

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

Number of fragments

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

Minimum acknowledgment overhead

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

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

Controlled latency

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

Optional congestion control

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

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

Appendix C. Considerations On Flow Control

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

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

Congestion on the forward path can also be indicated by an Explicit Congestion Notification (ECN) mechanism. Though whether and how ECN [RFC3168] is carried out over the LoWPAN is out of scope, this draft provides a way for the destination endpoint to echo an ECN indication back to the source endpoint in an acknowledgment message as represented in Figure 5 in Section 4.2.

It must be noted that congestion and collision are different topics. In particular, when a mesh operates on a same channel over multiple hops, then the forwarding of a fragment over a certain hop may collide with the forwarding of a next fragment that is following over a previous hop but in a same interference domain. This draft enables an end-to-end flow control, but leaves it to the sender stack to pace individual fragments within a transmit window, so that a given fragment is sent only when the previous fragment has had a chance to progress beyond the interference domain of this hop. In the case of 6TiSCH [I-D.ietf-6tisch-architecture], which operates over the TimeSlotted Channel Hopping [RFC7554] (TSCH) mode of operation of IEEE802.14.5, a fragment is forwarded over a different channel at a different time and it 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 5 describes how the sender decides how many fragments are (re)sent before an acknowledgment is required, and how the sender adapts that number to the network conditions.

Authors' Addresses

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

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

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

Email: jonhui@nestlabs.com

6lo
Internet-Draft
Intended status: Informational
Expires: September 7, 2018

T. Watteyne, Ed.
Analog Devices
C. Bormann
Universitaet Bremen TZI
P. Thubert
Cisco
March 06, 2018

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

Abstract

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

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 7, 2018.

Copyright Notice

Copyright (c) 2018 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. Critique of VRB	7
5. Security Considerations	8
6. IANA Considerations	8
7. Acknowledgments	8
8. Informative References	8
Authors' Addresses	8

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 header compression 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.

A reassembly buffer for 6LoWPAN contains:

- o datagram_size,
- o datagram_tag and link-layer sender and receiver addresses (to which the datagram_tag is local),
- o 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 the partial packet after a timeout.

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

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

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

2. Limits of Per-Hop Fragmentation and Reassembly

There are at least 2 limits to doing per-hop fragmentation and reassembly:

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 the case where each fragment would be forwarded without per-hop reassembly.

2.2. Memory Management and Reliability

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

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




Figure 2: Illustrating the Memory Management Issue.

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

3. Virtual Reassembly Buffer (VRB) Implementation

One implementation of 6LoWPAN fragmentation overcomes the limits listed in Section 2. The idea is for a node to immediately retransmit a fragment it receives, without fully reassembling the packet. This idea was introduced in Section 2.5.2 of [BOOK]. That is, a node may attempt to send out the data for a fragment in the form of a forwarded fragment, as soon as all necessary information for that is available.

Obviously, all fragments need to be sent with the same outgoing address (otherwise a full reassembly implementation would discard the fragments) and the same `datagram_tag`.

We use Figure 3 to illustrate VRB, and focus on the behavior of node E. With VRB, node E maintains a VRB table which functions similarly to a switching table: when receiving a fragment from node B with `datagram_tag=2`, forward it to node F with `datagram_tag=8`.




Figure 3: Illustrating VRB. #(5) and %(1) are fragments from packets coming from nodes A and C, with `datagram_tag` set to 5 and 1, respectively.

The VRB table is initially empty. An implementation might have for example pre-allocate memory for a VRB table with 4 entries (as in Figure 3), initially cleared.

When node E receives fragment 1 from node B with `datagram_tag=2`, it inspects the contents of the fragment and reads out the destination IPv6 address. When it is not destined to it, node E identifies the next hop to send this fragment to. It then creates an entry in the VRB table which contains 4 fields: (1) the link-layer address of the sender of the fragment it received, (2) the `datagram_tag` of the fragment it received, (3) the link-layer address of the next hop, (4) a `datagram_tag` for the fragments it will send. The latter `datagram_tag` must be locally unique.

Note that, if node E had multiple interfaces, the VRB table would also need additional column to identify the incoming and outgoing interface.

Any subsequent fragment that matches the "incoming" columns in the node's VRB table are immediately forwarded using the information in the "outgoing" columns. Note that, while this results in a behavior similar to link-layer switching, what is really happening is that the node has a virtual reassembly buffer. That is, it operates as if the packet were reassembled and fragmented, without ever actually holding a fully reassembled packet in memory.

Upon forwarding the last fragment of a packet, the VRB table entry can be cleared, and reused for a future packet. If the last fragment of a packet is dropped, the VRB table entry can be invalidated by timeout. Its timeout value is set to a maximum of 60 seconds as the reassembly timeout defined in [RFC4944].

A simple implementation may do away with any attempt to keep packet data in the virtual reassembly buffer. It then has to discard all non-first fragments for which a reassembly buffer is not already available (penalizing reordering, which however may be rare).

In case fragments can come out of order (a rare case, as all fragments of a packet are sent between the same neighbors), an implementation can use multiple the following two techniques. In case fragment 1 isn't received first, it can temporarily buffer fragments 2, 3, etc., until fragment 1 is received, and a next hop neighbor can be identified. Similarly, as the final fragment of the packet isn't necessarily received last, an implementation can maintain a bitmap of already forwarded fragments to know when all fragments have been forwarded (and the corresponding VRB entry can be cleared).

Note that the decision to do local processing of a packet needs to be taken with the first fragment - such packets of course do need to be fully reassembled (unless transport and application also can cope with fragments, which they rarely can in the presence of security).

It is possible for a network to be composed of some nodes that implement VRB, and others that don't. Nodes that do not implement VRB reassemble the packet.

[RFC6282] defines the header compression format for 6LoWPAN. One important impact of header compression is that the header is no longer of a fixed length. In particular, changes made by a forwarder may gain or lose the ability to use a more highly compressed variant, changing the length of the header in the packet.

If the change increases the size, the maximum frame size may be exceeded, leading to the need to re-fragment in the forwarder. This is less of a problem with full reassembly, but with virtual reassembly can lead to the need for sending an additional frame for each packet.

The well-known approach to minimize the probability of this need is for the original sender to put all slack in the frame sizes into the `_first_` packet, making this the smallest fragment and not the last one as would be done in a naive implementation. (This also has other consequences related to delivery probability, which are not discussed here.) This makes sure an additional fragment only needs to be sent if the header expansion during forwarding would have created an additional fragment with full reassembly as well.

4. Critique of VRB

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

There are, however, limits:

Non-zero Packet Drop Probability: Each VRB table entry can be 12 B (assuming 16-bit link-layer addresses). This is a footprint 2 orders of magnitude smaller compared to needing a 1280-byte reassembly buffer for each packet. Yet, the size of the VRB table necessarily remains finite. In the extreme case where a node is required to concurrently forward more packets than it has entries in its VRB table, packets are dropped.

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

No Per-Fragment Routing: All subsequent fragments follow the same sequence of hops from the source to the destination node as fragment 1.

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

If the limits are both present and not accepted by the application, future specifications may define new protocols to overcome these

limits. One example is [I-D.thubert-6lo-fragment-recovery] which defines a protocol which allows fragment recovery.

5. Security Considerations

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

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

6. IANA Considerations

No requests to IANA are made by this document.

7. Acknowledgments

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

8. Informative References

- [BOOK] Shelby, Z. and C. Bormann, "6LoWPAN", John Wiley & Sons, Ltd monograph, DOI 10.1002/9780470686218, November 2009.
- [I-D.thubert-6lo-fragment-recovery] Thubert, P., "6LoWPAN Selective Fragment Recovery", draft-thubert-6lo-fragment-recovery-00 (work in progress), February 2018.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<https://www.rfc-editor.org/info/rfc4944>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<https://www.rfc-editor.org/info/rfc6282>>.

Authors' Addresses

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

Email: thomas.wattheyne@analog.com

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

Email: cabo@tzi.org

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

Email: pthubert@cisco.com