

lpwan Working Group  
Internet-Draft  
Intended status: Informational  
Expires: May 1, 2018

D. Barthel  
Orange SA  
L. Toutain  
Institut MINES TELECOM; IMT Atlantique  
A. Kandasamy  
Acklio  
October 28, 2017

LPWAN Static Context Header Compression (SCHC) for ICMPv6  
draft-barthel-icmpv6-schc-00

Abstract

ICMPv6 is a companion protocol to IPv6. It defines messages that inform the source of IPv6 packets of errors during packet delivery. It also defines the Echo Request/Reply messages that are used for basic network troubleshooting (ping command). ICMPv6 messages are transported on IPv6.

This document describes how to adapt ICMPv6 to Low Power Wide Area Networks (LPWANs) by compressing ICMPv6/IPv6 headers and by protecting the LPWAN network and the Device from undesirable ICMPv6 traffic.

Status of This Memo

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

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

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

This Internet-Draft will expire on May 1, 2018.

Copyright Notice

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

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

## Table of Contents

1. Introduction . . . . .	2
2. Terminology . . . . .	3
3. Use cases . . . . .	3
4. Detailed behavior . . . . .	4
4.1. Device is the source of an ICMPv6 error message . . . . .	4
4.2. Device is the destination of an ICMPv6 error message . . . . .	4
4.2.1. ICMPv6 error message compression. . . . .	5
4.3. Device does a ping . . . . .	6
4.4. Device is ping'ed . . . . .	8
5. Traceroute . . . . .	9
6. Security considerations . . . . .	10
7. IANA Considerations . . . . .	10
8. References . . . . .	11
8.1. Normative References . . . . .	11
8.2. Informative References . . . . .	11
Authors' Addresses . . . . .	12

## 1. Introduction

ICMPv6 [RFC4443] is a companion protocol to IPv6 [RFC8200].

[RFC4443] defines a generic message format. This format is used for messages to be sent back to the source of an IPv6 packet to inform it about errors during packet delivery.

More specifically, [RFC4443] defines 4 error messages: Destination Unreachable, Packet Too Big, Time Exceeded and Parameter Problem.

[RFC4443] also defines the Echo Request and Echo Reply messages, which provide support for the ping application.

Other ICMPv6 messages are defined in other RFCs, such as an extended format of the same messages [RFC4884] and other messages used by the Neighbor Discovery Protocol [RFC4861].

This document focuses on using Static Context Header Compression (SCHC) to compress [RFC4443] messages that need to be transmitted over the LPWAN network, and on having the LPWAN gateway proxying the Device to save it the unwanted traffic.

LPWANs' salient characteristics are described in [I-D.ietf-lpwan-overview]

## 2. Terminology

This draft re-uses the Terminology defined in [I-D.ietf-lpwan-ipv6-static-context-hc].

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. Use cases

In the LPWAN architecture, we can distinguish the following cases:

- o the Device is the (purported) source of an ICMP error message, mainly in response to an incorrect incoming IPv6 message, or in response to a ping request. In this case, as much as possible, the core SCHC C/D should act as a proxy and originate the ICMP message, so that the Device and the LPWAN network are protected from this unwanted traffic.
- o the Device is the destination of the ICMP message, mainly in response to a packet sent by the Device to the network that generates an error. In this case, we want the ICMP message to reach the Device, and this document describes in section Section 4.2.1 what SCHC compression should be applied.
- o the Device is the originator of an Echo Request message, and therefore the destination of the Echo Reply message.
- o the Device is the destination of an Echo Request message, and therefore the purported source of an Echo Reply message.

These cases are further described in Section 4.

4. Detailed behavior

4.1. Device is the source of an ICMPv6 error message

As stated in [RFC4443], a node should generate an ICMPv6 message in response to an IPv6 packet that is malformed or which cannot be processed due to some incorrect field value.

The general intent of this document is to spare both the Device and the LPWAN network this un-necessary traffic. The incorrect packets should be caught at the core SCHC C/D and the ICMPv6 notification should be sent back from there.




Figure 1: Example of ICMPv6 error message sent back to the Internet

Figure 1 shows an example of an IPv6 packet trying to reach a Device. Let's assume that the port number used as destination port is not "known" (needs better definition) from the core SCHC C/D. Instead of sending the packet over the LPWAN and having this packet rejected by the Device, the core SCHC C/D issues an ICMPv6 error message "Destination Unreachable" (Type 1) with Code 1 ("Port Unreachable") on behalf of the Device.

TODO: This assumes that all ports that the Device listens to will be matched by a SCHC rule. Is this the basic assumption of SCHC that all packets that do not match a rule are rejected? If yes, why do have fragmentation also for uncompressed packets?

TODO: discuss the various Type/Code that are expected to be generated in response to various errors.

4.2. Device is the destination of an ICMPv6 error message

In this situation, we assume that a Device has been configured to send information to a server on the Internet. If this server becomes no longer accessible, an ICMPv6 message will be generated back

towards the Device by an intermediate router. This information can be useful to the Device, for example for reducing the reporting rate in case of periodic reporting of data. Therefore, we compress the ICMPv6 message using SCHC and forward it to the Device over the LPWAN.




Figure 2: Example of ICMPv6 error message sent back to the Device

Figure 2 illustrates this behavior. The ICMPv6 error message is compressed as described in Section 4.2.1 and forwarded over the LPWAN to the Device.

4.2.1. ICMPv6 error message compression.

The ICMPv6 error messages defined in [RFC4443] contain the fields shown in Figure 3.

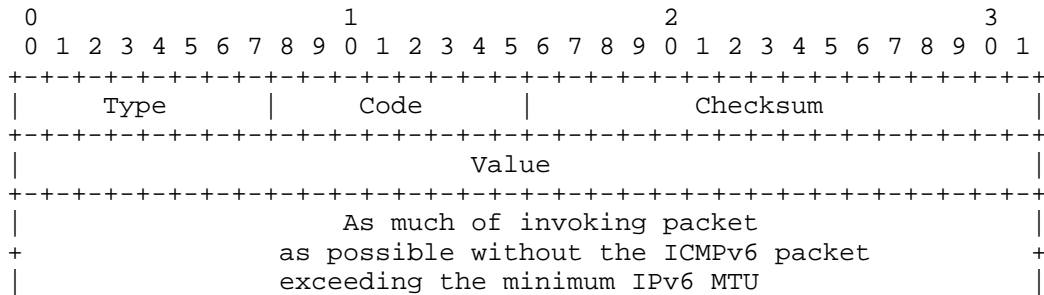


Figure 3: ICMPv6 Error Message format

[RFC4443] states that Type can take the values 1 to 4, and Code can be set to values between 0 and 6. Value is unused for the Destination Unreachable and Time Exceeded messages. It contains the MTU for the Packet Too Big message and a pointer to the byte causing

the error for the Parameter Error message. Therefore, Value is never expected to be greater than 1280 in LPWAN networks.

The following generic rule can therefore be used to compress all ICMPv6 error messages as defined today. More specific rules can also be defined to achieve better compression of some error messages.

The Type field can be associated to a matching list [1, 2, 3, 4] and is therefore compressed down to 2 bits. Code can be reduced to 3 bits using the LSB CDA. Value can be sent on 11 bits using the LSB CDA, but if the Device is known to send smaller packets, then the size of this field can be further reduced.

By [RFC4443], the rest of the ICMPv6 message must contain as much as possible of the IPv6 offending (invoking) packet that triggered this ICMPv6 error message. This information is used to try and identify the SCHC rule that was used to decompress the offending IPv6 packet. If the rule can be found then the Rule Id is added at the end of the compressed ICMPv6 message. Otherwise the compressed packet ends with the compressed Value field.

[RFC4443] states that the "ICMPv6 error message MUST include as much of the IPv6 offending (invoking) packet ... as possible". In order to comply with this requirement, if there is enough information in the incoming ICMPv6 message for the core SCHC C/D to identify the rule that has been used to decompress the erroneous IPv6 packet, this Rule Id must be sent in the compressed ICMPv6 message to the Device. TODO: the erroneous IPv6 packet header (not just the Rule Id) should be sent back. This includes the Rule Id and the compression residue. This means the SCHC C/D uses the context backwards (in the reverse direction). How does the Device know it must also use the context backwards?

TODO: how does one know that the "payload" of a compressed-header packet is in fact another compressed header?

#### 4.3. Device does a ping

If a ping request is generated by a Device, then SCHC compression applies.

The format of an ICMPv6 Echo Request message is described in Figure 4, with Type=128 and Code=0.




Figure 4: ICMPv6 Echo Request message format

If we assume that one rule will be devoted to compressing Echo Request messages, then Type and Code are known in the rule to be 128 and 0 and can therefore be elided with the not-sent CDA.

Checksum can be reconstructed with the compute-checksum CDA and therefore is not transmitted.

[RFC4443] states that Identifier and Sequence Number are meant to "aid in matching Echo Replies to this Echo Request" and that they "may be zero". Data is "zero or more bytes of arbitrary data".

We recommend that Identifier be zero, Sequence Number be a counter on 3 bits, and Data be zero bytes (absent). Therefore, Identifier is elided with the not-sent CDA, Sequence Number is transmitted on 3 bits with the LSB CDA and no Data is transmitted.

The transmission cost of the Echo Request message is therefore the size of the Rule Id + 3 bits.

When the destination receives the Echo Request message, it will respond back with a Echo Reply message. This message bears the same format as the Echo Request message but with Type = 129 (see Figure 4).

[RFC4443] states that the Identifier, Sequence Number and Data fields of the Echo Reply message shall contain the same values as the invoking Echo Request message. Therefore, a rule shall be used similar to that used for compressing the Echo Request message.

TODO: how about a shared rule for Echo Request and Echo Reply with an LSB(1) CDA on the Type field? Or exploiting the Up/Down direction field in the rule?

4.4. Device is ping'ed

If the Device is ping'ed (i.e., is the destination of an Echo Request message), the default behavior is to avoid propagating the Echo Request message over the LPWAN.

This is the recommended behavior with the Code 0 (default value) of the Echo Request message. In addition, this document defines two other Code values to achieve two other behaviors.

The resulting three behaviors are shown on Figure 5 and described below:




Figure 5: Examples of ICMPv6 Echo Request/Reply

- o Code = 0: The Echo Request message is not propagated on the LPWAN to the Device. If the SCHC C/D finds a rule in the context with the IPv6 address of the Device, it responds with an Echo Reply on behalf of the Device. If no rule is found with that IPv6 address, the SCHC C/D does not respond.

TODO: again, we are assuming that no compression rule is equivalent to the device not providing the service.



- o Code = 1: the SCHC C/D queries the NGW (or maintains a local database) and answers with the number of seconds since the Device last transmission.

TODO: what does it mean to answer "with the number of seconds ..."? There is no such field in an Echo Reply message. Do we overwrite one of the fields (Identifier, Sequencer Number, Data)? They are all supposed to be copied from the Echo Request. Do we change their definition with Code==1 or Code==2?

- o Code = 2: the SCHC C/D compresses the ICMPv6 message and forwards it to the Device. The Echo Reply message sent by the Device is also compressed. Since the Echo Request message comes from the Internet, the values of the Identifier, Sequence Number and Data fields cannot be known in advance, and therefore must be transmitted. However, it is likely that the Echo Request with Code 2 will be firewalled from the Internet and restricted to authorized users. Therefore, the Echo Request message can be assumed to have the same content as recommended in Section 4.3, and the same compression rules apply.

## 5. Traceroute

The `traceroute6` program sends successive probe packets destined to a chosen target but with the Hop Limit value successively incremented from the initial value 1.

It expects to receive a "Time Exceeded" (Type = 3) "Hop Limit" (Code = 0) ICMPv6 error message back from the successive routers along the path to the destination.

The probe packet is usually a UDP datagram, but can also be a TCP datagram or even an ICMPv6 message. The destination port is chosen in the unassigned range in hope that the destination, when eventually reached, will respond with a "Destination Unreachable" (Type = 1) "Port Unreachable" (Code = 4) ICMPv6 error message.

It is not anticipated that a Device will want to traceroute a destination on the Internet.

By contrast, a host on the Internet may attempt to traceroute an IPv6 address that is assigned to an LPWAN device. This is described in Figure 6.




Figure 6: Example of traceroute to the LPWAN Device

When the probe packet first reaches the core SCHC C/D, its remaining Hop Limit is 1. The core SCHC C/D will respond back with a "Time Exceeded" (Type = 3) "Hop Limit" (Code = 0) ICMPv6 error message. Later on, when the probe packet reaches the core SCHC C/D with a Hop Limit value of 2, the core SCHC C/D will, as explained in Section 4.1, answer back with a "Destination Unreachable" (Type = 1) "Port Unreachable" (Code = 4) ICMPv6 error message. This is what the traceroute6 command expects. Therefore, the traceroute6 command will work with LPWAN IPv6 destinations, except for the time displayed for the destination, which is actually the time to its proxy.

However, if the probe packet happens to hit a port that matches a SCHC rule for that Device, the packet will be compressed with this rule and sent over the LPWAN, which is unfortunate. Forwarding of packets to the Device over the LPWAN should only be done from authenticated/trusted sources anyway. Rate-limitation on top of authentication will mitigate this nuisance.

6. Security considerations

TODO

7. IANA Considerations

TODO

## 8. References

### 8.1. Normative References

- [I-D.ietf-lpwan-ipv6-static-context-hc]  
Minaburo, A., Toutain, L., and C. Gomez, "LPWAN Static Context Header Compression (SCHC) and fragmentation for IPv6 and UDP", draft-ietf-lpwan-ipv6-static-context-hc-07 (work in progress), October 2017.
- [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>>.
- [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>>.
- [RFC4884] Bonica, R., Gan, D., Tappan, D., and C. Pignataro, "Extended ICMP to Support Multi-Part Messages", RFC 4884, DOI 10.17487/RFC4884, April 2007, <<https://www.rfc-editor.org/info/rfc4884>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.

### 8.2. Informative References

- [I-D.ietf-lpwan-overview]  
Farrell, S., "LPWAN Overview", draft-ietf-lpwan-overview-07 (work in progress), October 2017.

Authors' Addresses

Dominique Barthel  
Orange SA  
28 chemin du Vieux Chene  
BP 98  
38243 Meylan Cedex  
France

Email: dominique.barthel@orange.com

Laurent Toutain  
Institut MINES TELECOM; IMT Atlantique  
2 rue de la Chataigneraie  
CS 17607  
35576 Cesson-Sevigne Cedex  
France

Email: laurent.toutain@imt-atlantique.fr

Arunprabhu Kandasamy  
Acklio  
2bis rue de la Chataigneraie  
35510 Cesson-Sevigne Cedex  
France

Email: arun@ackl.io

lpwan Working Group  
Internet-Draft  
Intended status: Informational  
Expires: January 3, 2019

A. Minaburo  
Acklio  
L. Toutain  
Institut MINES TELECOM; IMT Atlantique  
R. Andreasen  
Universidad de Buenos Aires  
July 02, 2018

LPWAN Static Context Header Compression (SCHC) for CoAP  
draft-ietf-lpwan-coap-static-context-hc-04

Abstract

This draft defines the way SCHC header compression can be applied to CoAP headers. CoAP header structure differs from IPv6 and UDP protocols since the CoAP use a flexible header with a variable number of options themselves of a variable length. Another important difference is the asymmetry in the header format used in request and response messages. Most of the compression mechanisms have been introduced in [I-D.ietf-lpwan-ipv6-static-context-hc], this document explains how to use the SCHC compression for CoAP.

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	3
2. SCHC Compression Process	3
3. CoAP Compression with SCHC	4
4. Compression of CoAP header fields	6
4.1. CoAP version field	6
4.2. CoAP type field	6
4.3. CoAP code field	6
4.4. CoAP Message ID field	6
4.5. CoAP Token fields	7
5. CoAP options	7
5.1. CoAP Content and Accept options.	7
5.2. CoAP option Max-Age field, CoAP option Uri-Host and Uri-Port fields	7
5.3. CoAP option Uri-Path and Uri-Query fields	8
5.3.1. Variable length Uri-Path and Uri-Query	8
5.3.2. Variable number of path or query elements	9
5.4. CoAP option Size1, Size2, Proxy-URI and Proxy-Scheme fields	9
5.5. CoAP option ETag, If-Match, If-None-Match, Location-Path and Location-Query fields	9
6. Other RFCs	9
6.1. Block	9
6.2. Observe	10
6.3. No-Response	10
6.4. Time Scale	10
6.5. OSCORE	10
7. Examples of CoAP header compression	12
7.1. Mandatory header with CON message	12
7.2. Complete exchange	13
7.3. OSCORE Compression	14
7.4. Example OSCORE Compression	17
8. Normative References	22
Authors' Addresses	22

## 1. Introduction

CoAP [rfc7252] is an implementation of the REST architecture for constrained devices. Nevertheless, if limited, the size of a CoAP header may be too large for LPWAN constraints and some compression may be needed to reduce the header size.

[I-D.ietf-lpwan-ipv6-static-context-hc] defines a header compression mechanism for LPWAN network based on a static context. The context is said static since the field description composing the Rules and the context are not learned during the packet exchanges but are previously defined. The context(s) is(are) known by both ends before transmission.

A context is composed of a set of rules that are referenced by Rule IDs (identifiers). A rule contains an ordered list of the fields descriptions containing a field ID (FID), its length (FL) and its position (FP), a direction indicator (DI) (upstream, downstream and bidirectional) and some associated Target Values (TV). Target Value indicates the value that can be expected. TV can also be a list of values. A Matching Operator (MO) is associated to each header field description. The rule is selected if all the MOs fit the TVs for all fields. In that case, a Compression/Decompression Action (CDA) associated to each field defines the link between the compressed and decompressed value for each of the header fields. Compression results mainly in 4 actions: send the field value, send nothing, send less significant bits of a field, send an index. Values sent are called Compression Residues and follows the rule ID.

## 2. SCHC Compression Process

The SCHC Compression rules can be applied to CoAP flows. SCHC Compression of the CoAP header may be done in conjunction with the above layers (IPv6/UDP) or independently. The SCHC adaptation layers as described in [I-D.ietf-lpwan-ipv6-static-context-hc] may be used as as shown in the Figure 1.




Figure 1: rule scope for CoAP

Figure 1 shows some examples for CoAP architecture and the SCHC rule's scope. A rule can covers all headers from IPv6 to CoAP, SCHC C/D is done in the device and at the LPWAN boundary. If an end-to-end encryption mechanisms is used between the device and the application. CoAP must be compressed independently of the other layers. The rule ID and the compression residue are encrypted using a mechanism such as DTLS. Only the other end can decipher the information.

Layers below may also be compressed using other SCHC rules (this is out of the scope of this document). OSCORE [I-D.ietf-core-object-security] can also define 2 rules to compress the CoAP message. A first rule focuses on the inner header and is end to end, a second rule may compress the outer header and the layer above. SCHC C/D for inner header is done by both ends, SCHC C/D for outer header and other headers is done between the device and the LPWAN boundary.

### 3. CoAP Compression with SCHC

CoAP differs from IPv6 and UDP protocols on the following aspects:

- o IPv6 and UDP are symmetrical protocols. The same fields are found in the request and in the response, only the location in the header may vary (e.g. source and destination fields). A CoAP request is different from a response. For example, the URI-path option is mandatory in the request and is not found in the response, a request may contain an Accept option and the response a Content option.

[I-D.ietf-lpwan-ipv6-static-context-hc] defines the use of a message direction (DI) when processing the rule which allows the description of message header format in both directions.



- o Even when a field is "symmetric" (i.e. found in both directions) the values carried in each direction are different. Combined with a matching list in the TV, this will allow to reduce the range of expected values in a particular direction and therefore reduce the size of a compression residue. For instance, if a client sends only CON request, the type can be elided by compression and the answer may use one bit to carry either the ACK or RST type. Same behavior can be applied to the CoAP Code field (0.OX code are present in the request and Y.ZZ in the answer). The direction allows to split in two parts the possible values for each direction.
- o In IPv6 and UDP header fields have a fixed size. In CoAP, Token size may vary from 0 to 8 bytes, length is given by a field in the header. More systematically, the CoAP options are described using the Type-Length-Value.

[I-D.ietf-lpwan-ipv6-static-context-hc] offers the possibility to define a function for the Field Length in the Field Description.

- o In CoAP headers, a field can be duplicated several times, for instances, elements of an URI (path or queries). The position defined in a rule, associated to a Field ID, can be used to identify the proper element.

[I-D.ietf-lpwan-ipv6-static-context-hc] allows a Field id to appears several times in the rule, the Field Position (FP) removes ambiguities for the matching operation.

- o Field size defined in the CoAP protocol can be too large regarding LPWAN traffic constraints. This is particularly true for the message ID field or Token field. The use of MSB MO can be used to reduce the information carried on LPWANs.
- o CoAP also obeys to the client/server paradigm and the compression rate can be different if the request is issued from an LPWAN node or from a non LPWAN device. For instance a Device (Dev) aware of LPWAN constraints can generate a 1 byte token, but a regular CoAP client will certainly send a larger token to the Thing. SCHC compression will not modify the values to offer a better compression rate. Nevertheless a proxy placed before the compressor may change some field values to offer a better compression rate and maintain the necessary context for interoperability with existing CoAP implementations.

#### 4. Compression of CoAP header fields

This section discusses of the compression of the different CoAP header fields.

##### 4.1. CoAP version field

This field is bidirectional and must be elided during the SCHC compression, since it always contains the same value. In the future, if new version of CoAP are defined, new rules ID will be defined avoiding ambiguities between versions.

##### 4.2. CoAP type field

[rfc7252] defines 4 types of messages: CON, NON, ACK and RST. The latter two ones are a response of the two first ones. If the device plays a specific role, a rule can exploit these property with the mapping list: [CON, NON] for one direction and [ACK, RST] for the other direction. Compression residue is reduced to 1 bit.

The field must be elided if for instance a client is sending only NON or CON messages.

In any case, a rule must be defined to carry RST to a client.

##### 4.3. CoAP code field

The compression of the CoAP code field follows the same principle as for the CoAP type field. If the device plays a specific role, the set of code values can be split in two parts, the request codes with the 0 class and the response values.

If the device implement only a CoAP client, the request code can be reduced to the set of request the client is able to process.

All the response codes should be compressed with a SCHC rule.

##### 4.4. CoAP Message ID field

This field is bidirectional and is used to manage acknowledgments. Server memorizes the value for a EXCHANGE\_LIFETIME period (by default 247 seconds) for CON messages and a NON\_LIFETIME period (by default 145 seconds) for NON messages. During that period, a server receiving the same Message ID value will process the message has a retransmission. After this period, it will be processed as a new messages.

In case the Device is a client, the size of the message ID field may be too large regarding the number of messages sent. Client may use only small message ID values, for instance 4 bit long. Therefore a MSB can be used to limit the size of the compression residue.

In case the Device is a server, client may be located outside of the LPWAN area and view the device as a regular device connected to the internet. The client will generate Message ID using the 16 bits space offered by this field. A CoAP proxy can be set before the SCHC C/D to reduce the value of the Message ID, to allow its compression with the MSB matching operator and LSB CDA.

#### 4.5. CoAP Token fields

Token is defined through two CoAP fields, Token Length in the mandatory header and Token Value directly following the mandatory CoAP header.

Token Length is processed as a tradition protocol field. If the value remains the same during all the transaction, the size can be stored in the context and elided during the transmission. Otherwise it will have to be sent as a compression residue.

Token Value size should not be defined directly in the rule in the Field Length (FL). Instead a specific function designed as "TKL" must be used. This function informs the SCHC C/D that the length of this field has to be read from the Token Length field.

### 5. CoAP options

#### 5.1. CoAP Content and Accept options.

These fields are both unidirectional and must not be set to bidirectional in a rule entry.

If single value is expected by the client, it can be stored in the TV and elided during the transmission. Otherwise, if several possible values are expected by the client, a matching-list should be used to limit the size of the residue. If not possible, the value has to be sent as a residue (fixed or variable length).

#### 5.2. CoAP option Max-Age field, CoAP option Uri-Host and Uri-Port fields

This field is unidirectional and must not be set to bidirectional in a rule entry. It is used only by the server to inform of the caching duration and is never found in client requests.

If the duration is known by both ends, value can be elided on the LPWAN.

A matching list can be used if some well-known values are defined.

Otherwise these options should be sent as a residue (fixed or variable length).

### 5.3. CoAP option Uri-Path and Uri-Query fields

This fields are unidirectional and must not be set to bidirectional in a rule entry. They are used only by the client to access to a specific resource and are never found in server responses.

Uri-Path and Uri-Query elements are a repeatable options, the Field Position (FP) gives the position in the path.

A Mapping list can be used to reduce size of variable Paths or Queries. In that case, to optimize the compression, several elements can be regrouped into a single entry. Numbering of elements do not change, MO comparison is set with the first element of the matching.

FID	FL	FP	DI	TV	MO	CDA
URI-Path	1	up		["/a/b", "/c/d"]	equal	not-sent
URI-Path	3	up			ignore	value-sent

Figure 2: complex path example

In Figure 2 a single bit residue can be used to code one of the 2 paths. If regrouping was not allowed, a 2 bits residue whould have been needed.

#### 5.3.1. Variable length Uri-Path and Uri-Query

When the length is known at the rule creation, the Field Length must be set to variable, and the unit is set to bytes.

The MSB MO can be apply to a Uri-Path or Uri-Query element. Since MSB value is given in bit, the size must always be a multiple of 8 bits and the LSB CDA must not carry any value.

The length sent at the beginning of a variable length residue indicates the size of the LSB in bytes.

For instance for a CoMi path /c/X6?k="eth0" the rule can be set to:

FID	FL	FP	DI	TV	MO	CDA
URI-Path	1	up	"c"	equal	not-sent	
URI-Path	2	up		ignore	value-sent	
URI-Query	1	up	"k="	MSB (16)	LSB	

Figure 3: CoMi URI compression

Figure 3 shows the parsing and the compression of the URI. where c is not sent. The second element is sent with the length (i.e. 0x2 X 6) followed by the query option (i.e. 0x05 "eth0").

### 5.3.2. Variable number of path or query elements

The number of Uri-path or Uri-Query element in a rule is fixed at the rule creation time. If the number varies, several rules should be created to cover all the possibilities. Another possibilities is to define the length of Uri-Path to variable and send a compression residue with a length of 0 to indicate that this Uri-Path is empty. This add 4 bits to the compression residue.

### 5.4. CoAP option Size1, Size2, Proxy-URI and Proxy-Scheme fields

These fields are unidirectional and must not be set to bidirectional in a rule entry. They are used only by the client to access to a specific resource and are never found in server response.

If the field value must be sent, TV is not set, MO is set to "ignore" and CDF is set to "value-sent". A mapping can also be used.

Otherwise the TV is set to the value, MO is set to "equal" and CDF is set to "not-sent"

### 5.5. CoAP option ETag, If-Match, If-None-Match, Location-Path and Location-Query fields

These fields are unidirectional.

These fields values cannot be stored in a rule entry. They must always be sent with the compression residues.

## 6. Other RFCs

### 6.1. Block

Block [rfc7959] allows a fragmentation at the CoAP level. SCHC includes also a fragmentation protocol. They are compatible. If a block option is used, its content must be sent as a compression residue.

## 6.2. Observe

[rfc7641] defines the Observe option. The TV is not set, MO is set to "ignore" and the CDF is set to "value-sent". SCHC does not limit the maximum size for this option (3 bytes). To reduce the transmission size either the device implementation should limit the value increase or a proxy can modify the incrementation.

Since RST message may be sent to inform a server that the client do not require Observe response, a rule must allow the transmission of this message.

## 6.3. No-Response

[rfc7967] defines an No-Response option limiting the responses made by a server to a request. If the value is not known by both ends, then TV is set to this value, MO is set to "equal" and CDF is set to "not-sent".

Otherwise, if the value is changing over time, TV is not set, MO is set to "ignore" and CDA to "value-sent". A matching list can also be used to reduce the size.

## 6.4. Time Scale

Time scale [I-D.toutain-core-time-scale] option allows a client to inform the server that it is in a slow network and that message ID should be kept for a duration given by the option.

If the value is not known by both ends, then TV is set to this value, MO is set to "equal" and CDA is set to "not-sent".

Otherwise, if the value is changing over time, TV is not set, MO is set to "ignore" and CDA to "value-sent". A matching list can also be used to reduce the size.

## 6.5. OSCORE

OSCORE [I-D.ietf-core-object-security] defines end-to-end protection for CoAP messages. This section describes how SCHC rules can be applied to compress OSCORE-protected messages.

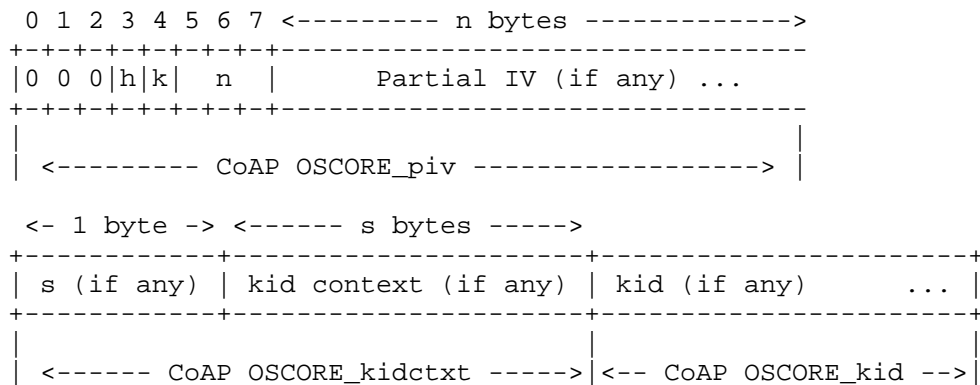


Figure 4: OSCORE Option

The encoding of the OSCORE Option Value defined in Section 6.1 of [I-D.ietf-core-object-security] is repeated in Figure 4.

The first byte is used for flags that specify the contents of the OSCORE option. The 3 most significant bits are reserved and always set to 0. Bit h, when set, indicates the presence of the kid context field in the option. Bit k, when set, indicates the presence of a kid field. The 3 least significant bits n indicate to length of the piv field in bytes, n = 0 taken to mean that no piv is present.

After the flag byte follow the piv field, kid context field and kid field in order and if present; the length of the kid context field is encoded in the first byte denoting by s the length of the kid context in bytes.

This draft recommends to implement a parser that is able to identify the OSCORE Option and the fields it contains - this makes it possible to do a preliminary processing of the message in preparation for regular SCHC compression.

Conceptually, the OSCORE option can transmit up to 3 distinct pieces of information at a time: the piv, the kid context, and the kid. This draft proposes that the SCHC Parser split the contents of this option into 3 SCHC fields:

- o CoAP OSCORE\_piv,
- o CoAP OSCORE\_ctxt,
- o CoAP OSCORE\_kid.

These fields are superposed on the OSCORE Option format in Figure 4, and include the corresponding flag and size bits for each part of the option. Both the flag and size bits can be omitted by use of the MSB matching operator on each field.

7. Examples of CoAP header compression

7.1. Mandatory header with CON message

In this first scenario, the LPWAN compressor receives from outside client a POST message, which is immediately acknowledged by the Device. For this simple scenario, the rules are described Figure 5.

Rule ID 1

Field	FL	FP	DI	Target Value	Match Opera.	CDA	Sent [bits]
CoAP version			bi	01	equal	not-sent	
CoAP version			bi	01	equal	not-sent	
CoAP Type			dw	CON	equal	not-sent	
CoAP Type			up	[ACK, RST]	match-map	matching-sent	T
CoAP TKL			bi	0	equal	not-sent	
CoAP Code			bi	ML1	match-map	matching-sent	CC CCC
CoAP MID			bi	0000	MSB(7 )	LSB(9)	M-ID
CoAP Uri-Path			dw	path	equal 1	not-sent	

Figure 5: CoAP Context to compress header without token

The version and Token Length fields are elided. Code has shrunk to 5 bits using a matching list. Uri-Path contains a single element indicated in the matching operator.

Figure 6 shows the time diagram of the exchange. A client in the Application Server sends a CON request. It can go through a proxy which reduces the message ID to a smallest value, with at least the 9 most significant bits equal to 0. SCHC Compression reduces the header sending only the Type, a mapped code and the least 9 significant bits of Message ID.



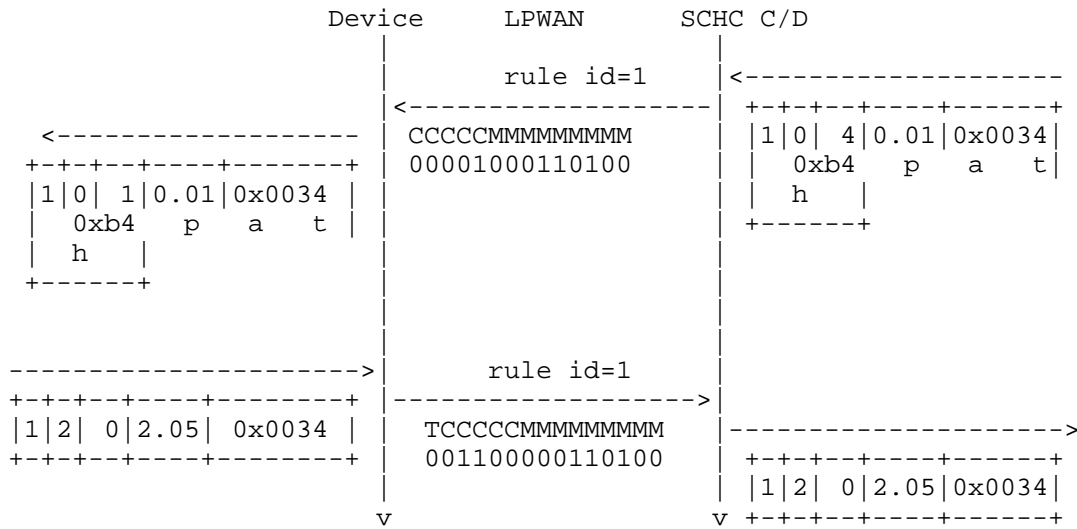



Figure 6: Compression with global addresses

7.2. Complete exchange

In that example, the Thing is using CoMi and sends queries for 2 SID.



### 7.3. OSCORE Compression

OSCORE aims to solve the problem of end-to-end encryption for CoAP messages, which are otherwise required to terminate their TLS or DTLS protection at the proxy, as discussed in Section 11.2 of [rfc7252]. CoAP proxies are men-in-the-middle, but not all of the information they have access to is necessary for their operation. The goal, therefore, is to hide as much of the message as possible while still enabling proxy operation.

Conceptually this is achieved by splitting the CoAP message into an Inner Plaintext and Outer OSCORE Message. The Inner Plaintext contains sensible information which is not necessary for proxy operation. This, in turn, is the part of the message which can be encrypted and need not be decrypted until it reaches its end destination. The Outer Message acts as a shell matching the format of a regular CoAP message, and includes all Options and information needed for proxy operation and caching. This decomposition is illustrated in Figure 7.

CoAP options are sorted into one of 3 classes, each granted a specific type of protection by the protocol:

- o Class E: Encrypted options moved to the Inner Plaintext,
- o Class I: Integrity-protected options included in the AAD for the encryption of the Plaintext but otherwise left untouched in the Outer Message,
- o Class U: Unprotected options left untouched in the Outer Message.

Additionally, the OSCORE Option is added as an Outer option, signaling that the message is OSCORE protected. This option carries the information necessary to retrieve the Security Context with which the message was encrypted so that it may be correctly decrypted at the other end-point.

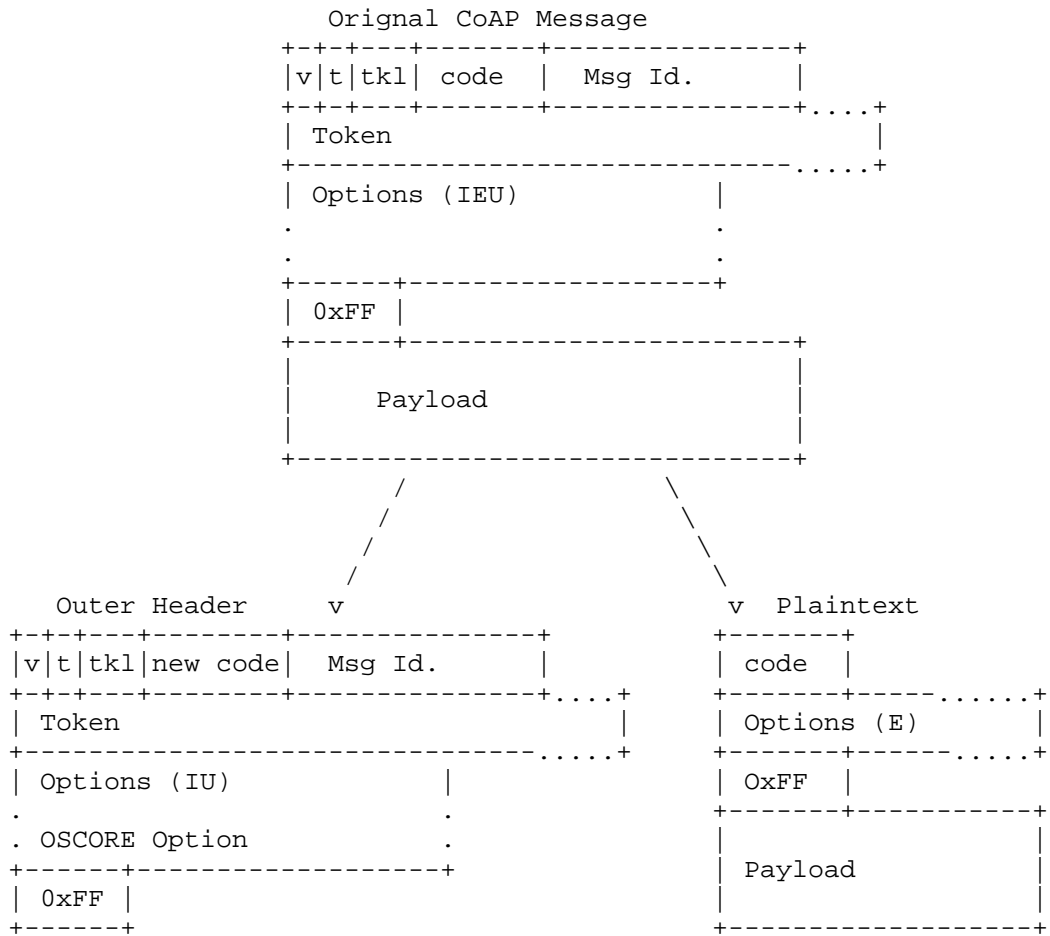


Figure 7: OSCORE inner and outer header form a CoAP message

Figure 7 shows the message format for the OSCORE Message and Plaintext. In the Outer Header, the original message code is hidden and replaced by a default value (POST or FETCH) depending on whether the original message was a Request or a Response. The original message code is put into the first byte of the Plaintext. Following the message code come the class E options and if present the original message Payload preceded by its payload marker.

The Plaintext is now encrypted by an AEAD algorithm which integrity protects Security Context parameters and eventually any class I options from the Outer Header. Currently no CoAP options are marked

class I. The resulting Ciphertext becomes the new Payload of the OSCORE message, as illustrated in Figure 8.

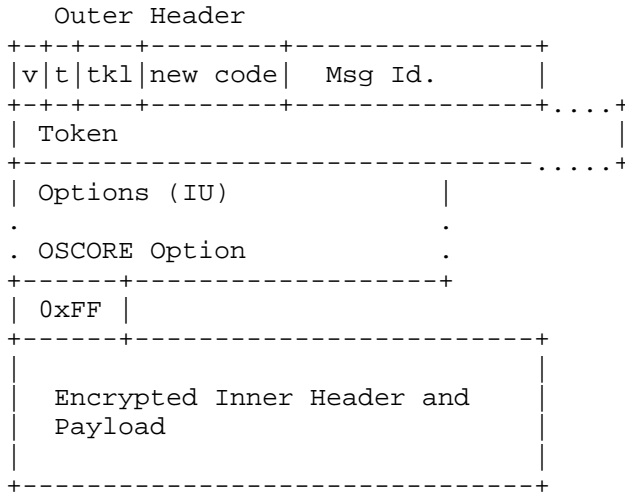


Figure 8: OSCORE message

The SCHC Compression scheme consists of compressing both the Plaintext before encryption and the resulting OSCORE message after encryption, see Figure 9. This way compression reaches all fields of the original CoAP message.

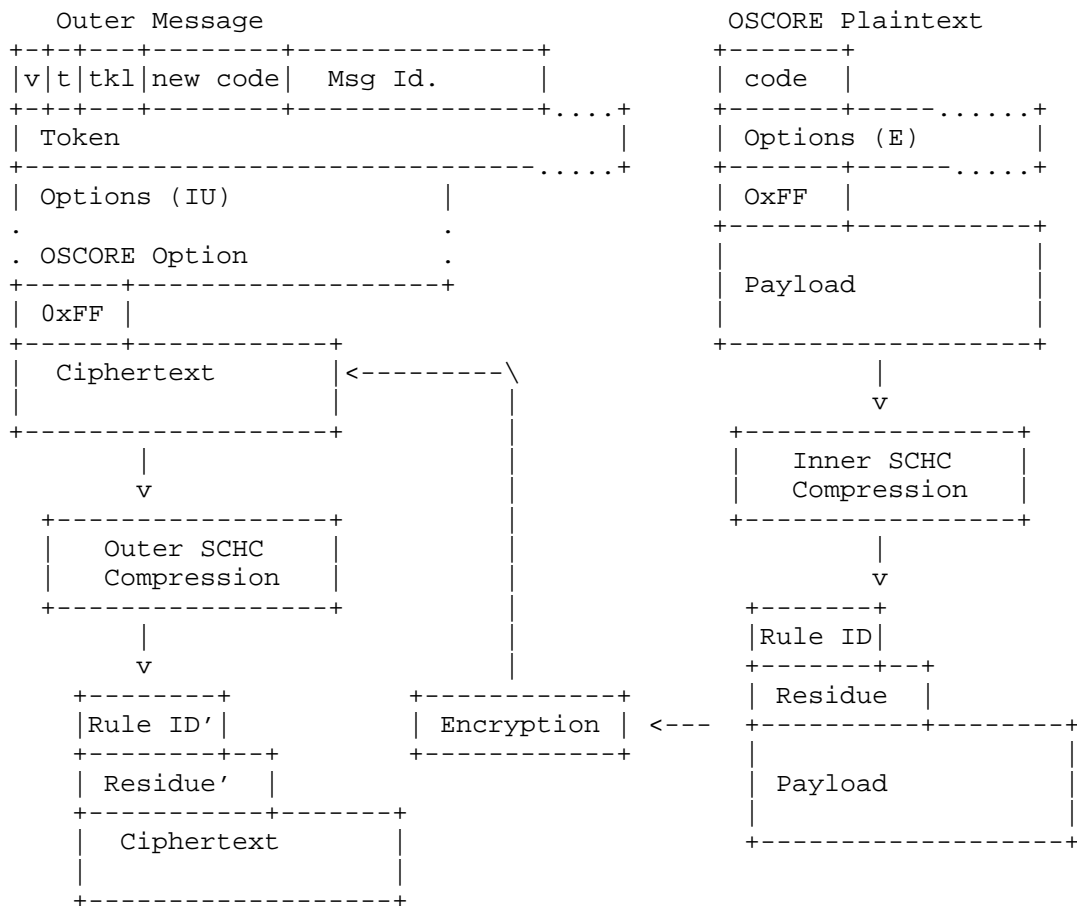


Figure 9: OSCORE Compression Diagram

#### 7.4. Example OSCORE Compression

In what follows we present an example GET Request and consequent CONTENT Response and show a possible set of rules for the Inner and Outer SCHC Compression. We then show a dump of the results and contrast SCHC + OSCORE performance with SCHC + COAP performance. This gives an approximation to the cost of security with SCHC-OSCORE.

Our first example CoAP message is the GET Request in Figure 10

```
Original message:
=====
0x4101000182bb74656d7065726174757265

Header:
0x4101
01 Ver
  00 CON
    0001 tk1
      00000001 Request Code 1 "GET"

0x0001 = mid
0x82 = token

Options:
0xbb74656d7065726174757265
Option 11: URI_PATH
Value = temperature
```

Original msg length: 17 bytes.

Figure 10: CoAP GET Request

Its corresponding response is the CONTENT Response in Figure 11.

```
Original message:
=====
0x6145000182ff32332043

Header:
0x6145
01 Ver
  10 ACK
    0001 tk1
      01000101 Successful Response Code 69 "2.05 Content"

0x0001 = mid
0x82 = token

0xFF Payload marker
Payload:
0x32332043
```

Original msg length: 10

Figure 11: CoAP CONTENT Response

The SCHC Rules for the Inner Compression include all fields that are already present in a regular CoAP message, what matters is the order of appearance and inclusion of only those CoAP fields that go into the Plaintext, Figure 12.

Rule ID 0

Field	FP	DI	Target Value	MO	CDA	Sent [bits]
CoAP Code		up	1	equal	not-sent	
CoAP Code		dw	[69,132]	match-map	match-sent	c
CoAP Uri-Path		up	temperature	equal	not-sent	
COAP Option-End		dw	0xFF	equal	not-sent	

Figure 12: Inner SCHC Rules

The Outer SCHC Rules (Figure 13) must process the OSCORE Options fields. Here we mask off the repeated bits (most importantly the flag and size bits) with the MSB Mathing Operator.

Rule ID 0

Field	FP	DI	Target Value	MO	CDA	Sent [bits]
CoAP version		bi	01	equal	not-sent	
CoAP Type		up	0	equal	not-sent	
CoAP Type		dw	2	equal	not-sent	
CoAP TKL		bi	1	equal	not-sent	
CoAP Code		up	2	equal	not-sent	
CoAP Code		dw	68	equal	not-sent	
CoAP MID		bi	0000	MSB(12)	LSB	MMMM
CoAP Token		bi	0x80	MSB(5)	LSB	TTT
CoAP OSCORE_piv		up	0x0900	MSB(12)	LSB	PPPP
COAP OSCORE_kid		up	b'\x06client'	MSB(52)	LSB	KKKK
CoAP OSCORE_piv		dw	b''	equal	not-sent	
COAP Option-End		dw	0xFF	equal	not-sent	

Figure 13: Outer SCHC Rules

Next we show a dump of the compressed message:

```
Compressed message:
=====
0x00291287f0a5c4833760d170
0x00 = Rule ID

piv = 0x04

Compression residue:
0b0001 010 0100 0100 (15 bits -> 2 bytes with padding)
  mid  tkn  piv  kid

Payload
0xafc297120cdd8345c

Compressed message length: 12 bytes
```

Figure 14: SCHC-OSCORE Compressed GET Request

```
Compressed message:
=====
0x0015f4de9cb814c96aed9b1d981a3a58
0x00 = Rule ID

Compression residue:
0b0001 010 (7 bits -> 1 byte with padding)
  mid  tkn

Payload
0xfa6f4e5c0a64b576cd8ecc0d1d2c

Compressed msg length: 16 bytes
```

Figure 15: SCHC-OSCORE Compressed CONTENT Response

For contrast, we compare these results with what would be obtained by SCHC compressing the original CoAP messages without protecting them with OSCORE. To do this, we compress the CoAP messages according to the SCHC rules in Figure 16.



Rule ID 1

Field	FP	DI	Target Value	MO	CDA	Sent [bits]
CoAP version		bi	01	equal	not-sent	
CoAP Type		up	0	equal	not-sent	
CoAP Type		dw	2	equal	not-sent	
CoAP TKL		bi	1	equal	not-sent	
CoAP Code		up	2	equal	not-sent	
CoAP Code		dw	[69,132]	equal	not-sent	
CoAP MID		bi	0000	MSB(12)	LSB	MMMM
CoAP Token		bi	0x80	MSB(5)	LSB	TTT
CoAP Uri-Path		up	temperature	equal	not-sent	
COAP Option-End		dw	0xFF	equal	not-sent	

Figure 16: SCHC-CoAP Rules (No OSCORE)

This yields the results in Figure 17 for the Request, and Figure 18 for the Response.

Compressed message:

```

=====
0x0114
0x01 = Rule ID
    
```

Compression residue:

```

0b00010100 (1 byte)
    
```

Compressed msg length: 2

Figure 17: CoAP GET Compressed without OSCORE

Compressed message:

```

=====
0x010a32332043
0x01 = Rule ID
    
```

Compression residue:

```

0b00001010 (1 byte)
    
```

Payload

```

0x32332043
    
```

Compressed msg length: 6

Figure 18: CoAP CONTENT Compressed without OSCORE

As can be seen, the difference between applying SCHC + OSCORE as compared to regular SCHC + CoAP is about 10 bytes of cost.

## 8. Normative References

[I-D.ietf-core-object-security]

Selander, G., Mattsson, J., Palombini, F., and L. Seitz, "Object Security for Constrained RESTful Environments (OSCORE)", draft-ietf-core-object-security-13 (work in progress), June 2018.

[I-D.ietf-lpwan-ipv6-static-context-hc]

Minaburo, A., Toutain, L., Gomez, C., and D. Barthel, "LPWAN Static Context Header Compression (SCHC) and fragmentation for IPv6 and UDP", draft-ietf-lpwan-ipv6-static-context-hc-16 (work in progress), June 2018.

[I-D.toutain-core-time-scale]

Minaburo, A. and L. Toutain, "CoAP Time Scale Option", draft-toutain-core-time-scale-00 (work in progress), October 2017.

[rfc7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/RFC7252, June 2014, <<https://www.rfc-editor.org/info/rfc7252>>.

[rfc7641] Hartke, K., "Observing Resources in the Constrained Application Protocol (CoAP)", RFC 7641, DOI 10.17487/RFC7641, September 2015, <<https://www.rfc-editor.org/info/rfc7641>>.

[rfc7959] Bormann, C. and Z. Shelby, Ed., "Block-Wise Transfers in the Constrained Application Protocol (CoAP)", RFC 7959, DOI 10.17487/RFC7959, August 2016, <<https://www.rfc-editor.org/info/rfc7959>>.

[rfc7967] Bhattacharyya, A., Bandyopadhyay, S., Pal, A., and T. Bose, "Constrained Application Protocol (CoAP) Option for No Server Response", RFC 7967, DOI 10.17487/RFC7967, August 2016, <<https://www.rfc-editor.org/info/rfc7967>>.

## Authors' Addresses

Ana Minaburo  
Acklio  
1137A avenue des Champs Blancs  
35510 Cesson-Sevigne Cedex  
France

Email: ana@ackl.io

Laurent Toutain  
Institut MINES TELECOM; IMT Atlantique  
2 rue de la Chataigneraie  
CS 17607  
35576 Cesson-Sevigne Cedex  
France

Email: Laurent.Toutain@imt-atlantique.fr

Ricardo Andreasen  
Universidad de Buenos Aires  
Av. Paseo Colon 850  
C1063ACV Ciudad Autonoma de Buenos Aires  
Argentina

Email: randreasen@fi.uba.ar

lpwan Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: December 31, 2018

A. Minaburo  
Acklio  
L. Toutain  
IMT-Atlantique  
C. Gomez  
Universitat Politecnica de Catalunya  
D. Barthel  
Orange Labs  
June 29, 2018

LPWAN Static Context Header Compression (SCHC) and fragmentation for  
IPv6 and UDP  
draft-ietf-lpwan-ipv6-static-context-hc-16

Abstract

This document defines the Static Context Header Compression (SCHC) framework, which provides both header compression and fragmentation functionalities. SCHC has been tailored for Low Power Wide Area Networks (LPWAN).

SCHC compression is based on a common static context stored in both the LPWAN devices and the network side. This document defines a header compression mechanism and its application to compress IPv6/UDP headers.

This document also specifies a fragmentation and reassembly mechanism that is used to support the IPv6 MTU requirement over the LPWAN technologies. Fragmentation is needed for IPv6 datagrams that, after SCHC compression or when such compression was not possible, still exceed the layer two maximum payload size.

The SCHC header compression and fragmentation mechanisms are independent of the specific LPWAN technology over which they are used. Note that this document defines generic functionalities and advisedly offers flexibility with regard to parameter settings and mechanism choices. Such settings and choices are expected to be made in other technology-specific documents.

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 31, 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 . . . . .	4
2. Requirements Notation . . . . .	5
3. LPWAN Architecture . . . . .	5
4. Terminology . . . . .	6
5. SCHC overview . . . . .	9
6. Rule ID . . . . .	13
7. Static Context Header Compression . . . . .	13
7.1. SCHC C/D Rules . . . . .	14
7.2. Rule ID for SCHC C/D . . . . .	16
7.3. Packet processing . . . . .	16
7.4. Matching operators . . . . .	18
7.5. Compression Decompression Actions (CDA) . . . . .	18
7.5.1. not-sent CDA . . . . .	20
7.5.2. value-sent CDA . . . . .	20
7.5.3. mapping-sent CDA . . . . .	20
7.5.4. LSB CDA . . . . .	20
7.5.5. DevIID, AppIID CDA . . . . .	21
7.5.6. Compute-* . . . . .	21
8. Fragmentation . . . . .	21
8.1. Overview . . . . .	21
8.2. Fragmentation Tools . . . . .	22

8.3.	Reliability modes . . . . .	25
8.4.	Fragmentation Formats . . . . .	27
8.4.1.	Fragments that are not the last one . . . . .	27
8.4.2.	All-1 fragment . . . . .	29
8.4.3.	SCHC ACK format . . . . .	31
8.4.4.	Abort formats . . . . .	33
8.5.	Baseline mechanism . . . . .	35
8.5.1.	No-ACK . . . . .	36
8.5.2.	ACK-Always . . . . .	36
8.5.3.	ACK-on-Error . . . . .	39
8.6.	Supporting multiple window sizes . . . . .	40
8.7.	Downlink SCHC Fragment transmission . . . . .	41
9.	Padding management . . . . .	42
10.	SCHC Compression for IPv6 and UDP headers . . . . .	43
10.1.	IPv6 version field . . . . .	43
10.2.	IPv6 Traffic class field . . . . .	43
10.3.	Flow label field . . . . .	44
10.4.	Payload Length field . . . . .	44
10.5.	Next Header field . . . . .	44
10.6.	Hop Limit field . . . . .	45
10.7.	IPv6 addresses fields . . . . .	45
10.7.1.	IPv6 source and destination prefixes . . . . .	45
10.7.2.	IPv6 source and destination IID . . . . .	46
10.8.	IPv6 extensions . . . . .	46
10.9.	UDP source and destination port . . . . .	46
10.10.	UDP length field . . . . .	47
10.11.	UDP Checksum field . . . . .	47
11.	IANA Considerations . . . . .	48
12.	Security considerations . . . . .	48
12.1.	Security considerations for SCHC Compression/Decompression . . . . .	48
12.2.	Security considerations for SCHC Fragmentation/Reassembly . . . . .	48
13.	Acknowledgements . . . . .	49
14.	References . . . . .	50
14.1.	Normative References . . . . .	50
14.2.	Informative References . . . . .	50
Appendix A.	SCHC Compression Examples . . . . .	51
Appendix B.	Fragmentation Examples . . . . .	54
Appendix C.	Fragmentation State Machines . . . . .	60
Appendix D.	SCHC Parameters - Ticket #15 . . . . .	67
Appendix E.	Note . . . . .	68
Authors' Addresses	. . . . .	69

## 1. Introduction

This document defines the Static Context Header Compression (SCHC) framework, which provides both header compression and fragmentation functionalities. SCHC has been tailored for Low Power Wide Area Networks (LPWAN).

Header compression is needed to efficiently bring Internet connectivity to the node within an LPWAN network. Some LPWAN networks properties can be exploited to get an efficient header compression:

- o The network topology is star-oriented, which means that all packets follow the same path. For the needs of this document, the architecture can simply be described as Devices (Dev) exchanging information with LPWAN Application Servers (App) through Network Gateways (NGW).
- o Because devices embed built-in applications, the traffic flows to be compressed are known in advance. Indeed, new applications cannot be easily installed in LPWAN devices, as they would in computers or smartphones.

The Static Context Header Compression (SCHC) is defined for this environment. SCHC uses a context, in which information about header fields is stored. This context is static: the values of the header fields do not change over time. This avoids complex resynchronization mechanisms, that would be incompatible with LPWAN characteristics. In most cases, a small context identifier is enough to represent the full IPv6/UDP headers. The SCHC header compression mechanism is independent of the specific LPWAN technology over which it is used.

LPWAN technologies impose some strict limitations on traffic. For instance, devices are sleeping most of the time and MAY receive data during short periods of time after transmission to preserve battery. LPWAN technologies are also characterized, among others, by a very reduced data unit and/or payload size (see [RFC8376]). However, some of these technologies do not provide fragmentation functionality, therefore the only option for them to support the IPv6 MTU requirement of 1280 bytes [RFC8200] is to use a fragmentation protocol at the adaptation layer, below IPv6. In response to this need, this document also defines a fragmentation/reassembly mechanism, which supports the IPv6 MTU requirement over LPWAN technologies. Such functionality has been designed under the assumption that there is no out-of-sequence delivery of data units between the entity performing fragmentation and the entity performing reassembly.

Note that this document defines generic functionality and purposefully offers flexibility with regard to parameter settings and mechanism choices. Such settings and choices are expected to be made in other, technology-specific documents.

2. Requirements Notation

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. LPWAN Architecture

LPWAN technologies have similar network architectures but different terminologies. Using the terminology defined in [RFC8376], we can identify different types of entities in a typical LPWAN network, see Figure 1:

- o Devices (Dev) are the end-devices or hosts (e.g. sensors, actuators, etc.). There can be a very high density of devices per radio gateway.
- o The Radio Gateway (RGW), which is the end point of the constrained link.
- o The Network Gateway (NGW) is the interconnection node between the Radio Gateway and the Internet.
- o LPWAN-AAA Server, which controls the user authentication and the applications.
- o Application Server (App)




Figure 1: LPWAN Architecture



#### 4. Terminology

This section defines the terminology and acronyms used in this document.

Note that the SCHC acronym is pronounced like "sheek" in English (or "chic" in French). Therefore, this document writes "a SCHC Packet" instead of "an SCHC Packet".

- o Abort. A SCHC Fragment format to signal the other end-point that the on-going fragment transmission is stopped and finished.
- o All-0. The SCHC Fragment format for the last fragment of a window that is not the last one of a SCHC Packet (see window in this glossary).
- o All-1. The SCHC Fragment format for the last fragment of the SCHC Packet.
- o All-0 empty. An All-0 SCHC Fragment without payload. It is used to request the SCHC ACK with the encoded Bitmap when the Retransmission Timer expires, in a window that is not the last one of a packet.
- o All-1 empty. An All-1 SCHC Fragment without payload. It is used to request the SCHC ACK with the encoded Bitmap when the Retransmission Timer expires in the last window of a packet.
- o App: LPWAN Application. An application sending/receiving IPv6 packets to/from the Device.
- o AppIID: Application Interface Identifier. The IID that identifies the application server interface.
- o Bi: Bidirectional. Characterises a Rule Entry that applies to headers of packets travelling in either direction (Up and Dw, see this glossary).
- o Bitmap: a bit field in the SCHC ACK message that tells the sender which SCHC Fragments in a window of fragments were correctly received.
- o C: Checked bit. Used in an acknowledgement (SCHC ACK) header to determine if the MIC locally computed by the receiver matches (1) the received MIC or not (0).
- o CDA: Compression/Decompression Action. Describes the reciprocal pair of actions that are performed at the compressor to compress a

header field and at the decompressor to recover the original header field value.

- o Compression Residue. The bits that need to be sent (beyond the Rule ID itself) after applying the SCHC compression over each header field.
- o Context: A set of Rules used to compress/decompress headers.
- o Dev: Device. A node connected to an LPWAN. A Dev SHOULD implement SCHC.
- o DevIID: Device Interface Identifier. The IID that identifies the Dev interface.
- o DI: Direction Indicator. This field tells which direction of packet travel (Up, Dw or Bi) a Rule applies to. This allows for asymmetric processing.
- o DTag: Datagram Tag. This SCHC F/R header field is set to the same value for all SCHC Fragments carrying the same SCHC Packet.
- o Dw: Downlink direction for compression/decompression in both sides, from SCHC C/D in the network to SCHC C/D in the Dev.
- o FCN: Fragment Compressed Number. This SCHC F/R header field carries an efficient representation of a larger-sized fragment number.
- o Field Description. A line in the Rule table.
- o FID: Field Identifier. This is an index to describe the header fields in a Rule.
- o FL: Field Length is the length of the packet header field. It is expressed in bits for header fields of fixed lengths or as a type (e.g. variable, token length, ...) for field lengths that are unknown at the time of Rule creation. The length of a header field is defined in the corresponding protocol specification.
- o FP: Field Position is a value that is used to identify the position where each instance of a field appears in the header.
- o IID: Interface Identifier. See the IPv6 addressing architecture [RFC7136]

- o Inactivity Timer. A timer used after receiving a SCHC Fragment to detect when, due to a communication error, there is no possibility to continue an on-going fragmented SCHC Packet transmission.
- o L2: Layer two. The immediate lower layer SCHC interfaces with. It is provided by an underlying LPWAN technology.
- o L2 Word: this is the minimum subdivision of payload data that the L2 will carry. In most L2 technologies, the L2 Word is an octet. In bit-oriented radio technologies, the L2 Word might be a single bit. The L2 Word size is assumed to be constant over time for each device.
- o MIC: Message Integrity Check. A SCHC F/R header field computed over the fragmented SCHC Packet and potential fragment padding, used for error detection after SCHC Packet reassembly.
- o MO: Matching Operator. An operator used to match a value contained in a header field with a value contained in a Rule.
- o Padding (P). Extra bits that may be appended by SCHC to a data unit that it passes to the underlying Layer 2 for transmission. SCHC itself operates on bits, not bytes, and does not have any alignment prerequisite. See Section 9.
- o Retransmission Timer. A timer used by the SCHC Fragment sender during an on-going fragmented SCHC Packet transmission to detect possible link errors when waiting for a possible incoming SCHC ACK.
- o Rule: A set of header field values.
- o Rule entry: A column in a Rule that describes a parameter of the header field.
- o Rule ID: An identifier for a Rule. SCHC C/D on both sides share the same Rule ID for a given packet. A set of Rule IDs are used to support SCHC F/R functionality.
- o SCHC ACK: A SCHC acknowledgement for fragmentation. This message is used to report on the success of reception of a set of SCHC Fragments. See Section 8 for more details.
- o SCHC C/D: Static Context Header Compression Compressor/Decompressor. A mechanism used on both sides, at the Dev and at the network, to achieve Compression/Decompression of headers. SCHC C/D uses Rules to perform compression and decompression.

- o SCHC F/R: Static Context Header Compression Fragmentation/Reassembly. A protocol used on both sides, at the Dev and at the network, to achieve Fragmentation/Reassembly of SCHC Packets. SCHC F/R has three reliability modes.
- o SCHC Fragment: A data unit that carries a subset of a SCHC Packet. SCHC F/R is needed when the size of a SCHC packet exceeds the available payload size of the underlying L2 technology data unit. See Section 8.
- o SCHC Packet: A packet (e.g. an IPv6 packet) whose header has been compressed as per the header compression mechanism defined in this document. If the header compression process is unable to actually compress the packet header, the packet with the uncompressed header is still called a SCHC Packet (in this case, a Rule ID is used to indicate that the packet header has not been compressed). See Section 7 for more details.
- o TV: Target value. A value contained in a Rule that will be matched with the value of a header field.
- o Up: Uplink direction for compression/decompression in both sides, from the Dev SCHC C/D to the network SCHC C/D.
- o W: Window bit. A SCHC Fragment header field used in ACK-on-Error or ACK-Always mode Section 8, which carries the same value for all SCHC Fragments of a window.
- o Window: A subset of the SCHC Fragments needed to carry a SCHC Packet (see Section 8).

## 5. SCHC overview

SCHC can be abstracted as an adaptation layer between IPv6 and the underlying LPWAN technology. SCHC comprises two sublayers (i.e. the Compression sublayer and the Fragmentation sublayer), as shown in Figure 2.


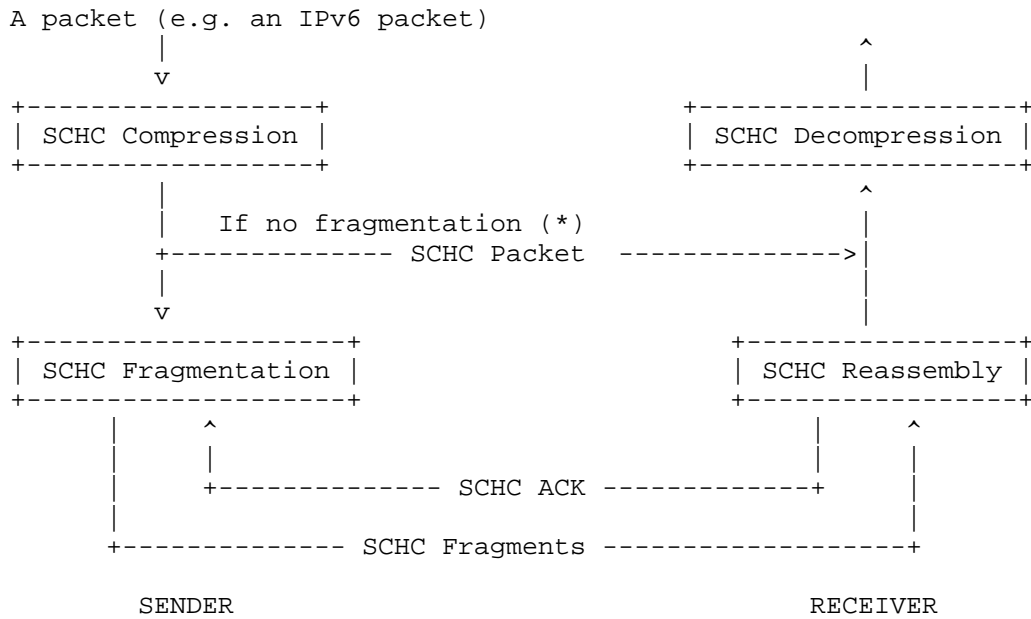


Figure 2: Protocol stack comprising IPv6, SCHC and an LPWAN technology

As per this document, when a packet (e.g. an IPv6 packet) needs to be transmitted, header compression is first applied to the packet. The resulting packet after header compression (whose header may or may not actually be smaller than that of the original packet) is called a SCHC Packet. If the SCHC Packet size exceeds the layer 2 (L2) MTU, fragmentation is then applied to the SCHC Packet. The SCHC Packet or the SCHC Fragments are then transmitted over the LPWAN. The reciprocal operations take place at the receiver. This process is illustrated in Figure 3.



\*: the decision to use Fragmentation or not is left to each LPWAN technology over which SCHC is applied. See LPWAN technology-specific documents.

Figure 3: SCHC operations taking place at the sender and the receiver

The SCHC Packet is composed of the Compressed Header followed by the payload from the original packet (see Figure 4). The Compressed Header itself is composed of a Rule ID and a Compression Residue. The Compression Residue may be absent, see Section 7. Both the Rule ID and the Compression Residue potentially have a variable size, and generally are not a multiple of bytes in size.




Figure 4: SCHC Packet

The Fragment Header size is variable and depends on the Fragmentation parameters. The Fragment payload contains a part of the SCHC Packet Compressed Header, a part of the SCHC Packet Payload or both. Its

size depends on the L2 data unit, see Section 8. The SCHC Fragment has the following format:




Figure 5: SCHC Fragment

The SCHC ACK is only used for Fragmentation. It has the following format:




Figure 6: SCHC ACK

The SCHC ACK Header and the encoded Bitmap both have variable size.

Figure 7 below maps the functional elements of Figure 3 onto the LPWAN architecture elements of Figure 1.




Figure 7: Architecture

SCHC C/D and SCHC F/R are located on both sides of the LPWAN transmission, i.e. on the Dev side and on the Network side.

Let's describe the operation in the Uplink direction. The Device application packets use IPv6 or IPv6/UDP protocols. Before sending

these packets, the Dev compresses their headers using SCHC C/D and, if the SCHC Packet resulting from the compression exceeds the maximum payload size of the underlying LPWAN technology, SCHC F/R is performed (see Section 8). The resulting SCHC Fragments are sent as one or more L2 frames to an LPWAN Radio Gateway (RG) which forwards them to a Network Gateway (NGW). The NGW sends the data to a SCHC F/R and then to the SCHC C/D for decompression. The SCHC F/R and C/D on the Network side can be located in the NGW or somewhere else as long as a tunnel is established between them and the NGW. Note that, for some LPWAN technologies, it MAY be suitable to locate the SCHC F/R functionality nearer the NGW, in order to better deal with time constraints of such technologies. The SCHC C/D and F/R on both sides MUST share the same set of Rules. After decompression, the packet can be sent over the Internet to one or several LPWAN Application Servers (App).

The SCHC C/D and F/R process is symmetrical, therefore the description of the Downlink direction trivially derives from the one above.

## 6. Rule ID

Rule IDs are identifiers used to select the correct context either for Compression/Decompression or for Fragmentation/Reassembly.

The size of the Rule IDs is not specified in this document, as it is implementation-specific and can vary according to the LPWAN technology and the number of Rules, among others.

The Rule IDs are used:

- o In the SCHC C/D context, to identify the Rule (i.e., the set of Field Descriptions) that is used to compress a packet header.
- o At least one Rule ID MAY be allocated to tagging packets for which SCHC compression was not possible (no matching Rule was found).
- o In SCHC F/R, to identify the specific modes and settings of SCHC Fragments being transmitted, and to identify the SCK ACKs, including their modes and settings. Note that in the case of bidirectional communication, at least two Rule ID values are therefore needed for F/R.

## 7. Static Context Header Compression

In order to perform header compression, this document defines a mechanism called Static Context Header Compression (SCHC), which is based on using context, i.e. a set of Rules to compress or decompress



headers. SCHC avoids context synchronization, which is the most bandwidth-consuming operation in other header compression mechanisms such as RoHC [RFC5795]. Since the nature of packets is highly predictable in LPWAN networks, static contexts MAY be stored beforehand to omit transmitting some information over the air. The contexts MUST be stored at both ends, and they can be learned by a provisioning protocol or by out of band means, or they can be pre-provisioned. The way the contexts are provisioned on both ends is out of the scope of this document.

7.1. SCHC C/D Rules

The main idea of the SCHC compression scheme is to transmit the Rule ID to the other end instead of sending known field values. This Rule ID identifies a Rule that provides the closest match to the original packet values. Hence, when a value is known by both ends, it is only necessary to send the corresponding Rule ID over the LPWAN network. How Rules are generated is out of the scope of this document. The Rules MAY be changed at run-time but the way to do this will be specified in another document.

The context contains a list of Rules (cf. Figure 8). Each Rule itself contains a list of Field Descriptions composed of a Field Identifier (FID), a Field Length (FL), a Field Position (FP), a Direction Indicator (DI), a Target Value (TV), a Matching Operator (MO) and a Compression/Decompression Action (CDA).




Figure 8: Compression/Decompression Context

A Rule does not describe how to parse a packet header to find each field. This MUST be known from the compressor/decompressor. Rules only describe the compression/decompression behavior for each header field. In a Rule, the Field Descriptions are listed in the order in which the fields appear in the packet header.

A Rule also describes what Compression Residue is sent. The Compression Residue is assembled by concatenating the residues for each field, in the order the Field Descriptions appear in the Rule.

The Context describes the header fields and its values with the following entries:

- o Field ID (FID) is a unique value to define the header field.
- o Field Length (FL) represents the length of the field. It can be either a fixed value (in bits) if the length is known when the Rule is created or a type if the length is variable. The length of a header field is defined in the corresponding protocol specification. The type defines the process to compute the length, its unit (bits, bytes,...) and the value to be sent before the Compression Residue.
- o Field Position (FP): most often, a field only occurs once in a packet header. Some fields may occur multiple times in a header. FP indicates which occurrence this Field Description applies to. The default value is 1 (first occurrence).
- o A Direction Indicator (DI) indicates the packet direction(s) this Field Description applies to. Three values are possible:
  - \* UPLINK (Up): this Field Description is only applicable to packets sent by the Dev to the App,
  - \* DOWNLINK (Dw): this Field Description is only applicable to packets sent from the App to the Dev,
  - \* BIDIRECTIONAL (Bi): this Field Description is applicable to packets travelling both Up and Dw.
- o Target Value (TV) is the value used to make the match with the packet header field. The Target Value can be of any type (integer, strings, etc.). For instance, it can be a single value or a more complex structure (array, list, etc.), such as a JSON or a CBOR structure.
- o Matching Operator (MO) is the operator used to match the Field Value and the Target Value. The Matching Operator may require

some parameters. MO is only used during the compression phase. The set of MOs defined in this document can be found in Section 7.4.

- o Compression Decompression Action (CDA) describes the compression and decompression processes to be performed after the MO is applied. Some CDAs MAY require parameter values for their operation. CDAs are used in both the compression and the decompression functions. The set of CDAs defined in this document can be found in Section 7.5.

### 7.2. Rule ID for SCHC C/D

Rule IDs are sent by the compression function in one side and are received for the decompression function in the other side. In SCHC C/D, the Rule IDs are specific to a Dev. Hence, multiple Dev instances MAY use the same Rule ID to define different header compression contexts. To identify the correct Rule ID, the SCHC C/D needs to correlate the Rule ID with the Dev identifier to find the appropriate Rule to be applied.

### 7.3. Packet processing

The compression/decompression process follows several steps:

- o Compression Rule selection: The goal is to identify which Rule(s) will be used to compress the packet's headers. When doing decompression, on the network side the SCHC C/D needs to find the correct Rule based on the L2 address and in this way, it can use the DevIID and the Rule ID. On the Dev side, only the Rule ID is needed to identify the correct Rule since the Dev only holds Rules that apply to itself. The Rule will be selected by matching the Fields Descriptions to the packet header as described below. When the selection of a Rule is done, this Rule is used to compress the header. The detailed steps for compression Rule selection are the following:
  - \* The first step is to choose the Field Descriptions by their direction, using the Direction Indicator (DI). A Field Description that does not correspond to the appropriate DI will be ignored. If all the fields of the packet do not have a Field Description with the correct DI, the Rule is discarded and SCHC C/D proceeds to explore the next Rule.
  - \* When the DI has matched, then the next step is to identify the fields according to Field Position (FP). If FP does not correspond, the Rule is not used and the SCHC C/D proceeds to consider the next Rule.

- \* Once the DI and the FP correspond to the header information, each packet field's value is then compared to the corresponding Target Value (TV) stored in the Rule for that specific field using the matching operator (MO).

If all the fields in the packet's header satisfy all the matching operators (MO) of a Rule (i.e. all MO results are True), the fields of the header are then compressed according to the Compression/Decompression Actions (CDAs) and a compressed header (with possibly a Compression Residue) SHOULD be obtained. Otherwise, the next Rule is tested.

- \* If no eligible Rule is found, then the header MUST be sent without compression. This MAY require the use of the SCHC F/R process.
- o Sending: If an eligible Rule is found, the Rule ID is sent to the other end followed by the Compression Residue (which could be empty) and directly followed by the payload. The Compression Residue is the concatenation of the Compression Residues for each field according to the CDAs for that Rule. The way the Rule ID is sent depends on the specific underlying LPWAN technology. For example, it can be either included in an L2 header or sent in the first byte of the L2 payload. (Cf. Figure 9). This process will be specified in the LPWAN technology-specific document and is out of the scope of the present document. On LPWAN technologies that are byte-oriented, the compressed header concatenated with the original packet payload is padded to a multiple of 8 bits, if needed. See Section 9 for details.
- o Decompression: When doing decompression, on the network side the SCHC C/D needs to find the correct Rule based on the L2 address and in this way, it can use the DevIID and the Rule ID. On the Dev side, only the Rule ID is needed to identify the correct Rule since the Dev only holds Rules that apply to itself.

The receiver identifies the sender through its device-id (e.g. MAC address, if exists) and selects the appropriate Rule from the Rule ID. If a source identifier is present in the L2 technology, it is used to select the Rule ID. This Rule describes the compressed header format and associates the values to the header fields. The receiver applies the CDA action to reconstruct the original header fields. The CDA application order can be different from the order given by the Rule. For instance, Compute-\* SHOULD be applied at the end, after all the other CDAs.




Figure 9: SCHC C/D Packet Format

#### 7.4. Matching operators

Matching Operators (MOs) are functions used by both SCHC C/D endpoints involved in the header compression/decompression. They are not typed and can be indifferently applied to integer, string or any other data type. The result of the operation can either be True or False. MOs are defined as follows:

- o equal: The match result is True if a field value in a packet and the value in the TV are equal.
- o ignore: No check is done between a field value in a packet and a TV in the Rule. The result of the matching is always true.
- o MSB(x): A match is obtained if the most significant x bits of the packet header field value are equal to the TV in the Rule. The x parameter of the MSB MO indicates how many bits are involved in the comparison. If the FL is described as variable, the length must be a multiple of the unit. For example, x must be multiple of 8 if the unit of the variable length is in bytes.
- o match-mapping: With match-mapping, the Target Value is a list of values. Each value of the list is identified by a short ID (or index). Compression is achieved by sending the index instead of the original header field value. This operator matches if the header field value is equal to one of the values in the target list.

#### 7.5. Compression Decompression Actions (CDA)

The Compression Decompression Action (CDA) describes the actions taken during the compression of headers fields, and inversely, the action taken by the decompressor to restore the original value.

Action	Compression	Decompression
not-sent	elided	use value stored in context
value-sent	send	build from received value
mapping-sent	send index	value from index on a table
LSB	send LSB	TV, received value
compute-length	elided	compute length
compute-checksum	elided	compute UDP checksum
DevIID	elided	build IID from L2 Dev addr
AppIID	elided	build IID from L2 App addr

Figure 10: Compression and Decompression Actions

Figure 10 summarizes the basic functions that can be used to compress and decompress a field. The first column lists the actions name. The second and third columns outline the reciprocal compression/decompression behavior for each action.

Compression is done in order that Fields Descriptions appear in a Rule. The result of each Compression/Decompression Action is appended to the working Compression Residue in that same order. The receiver knows the size of each compressed field which can be given by the Rule or MAY be sent with the compressed header.

If the field is identified as being variable in the Field Description, then the size of the Compression Residue value (using the unit defined in the FL) MUST be sent first using the following coding:

- o If the size is between 0 and 14, it is sent as a 4-bits integer.
- o For values between 15 and 254, the first 4 bits sent are set to 1 and the size is sent using 8 bits integer.
- o For higher values of size, the first 12 bits are set to 1 and the next two bytes contain the size value as a 16 bits integer.

If a field is not present in the packet but exists in the Rule and its FL is specified as being variable, size 0 MUST be sent to denote its absence.

#### 7.5.1. not-sent CDA

The not-sent function is generally used when the field value is specified in a Rule and therefore known by both the Compressor and the Decompressor. This action is generally used with the "equal" MO. If MO is "ignore", there is a risk to have a decompressed field value different from the original field that was compressed.

The compressor does not send any Compression Residue for a field on which not-sent compression is applied.

The decompressor restores the field value with the Target Value stored in the matched Rule identified by the received Rule ID.

#### 7.5.2. value-sent CDA

The value-sent action is generally used when the field value is not known by both the Compressor and the Decompressor. The value is sent as a residue in the compressed message header. Both Compressor and Decompressor MUST know the size of the field, either implicitly (the size is known by both sides) or by explicitly indicating the length in the Compression Residue, as defined in Section 7.5. This function is generally used with the "ignore" MO.

#### 7.5.3. mapping-sent CDA

The mapping-sent is used to send a smaller index (the index into the Target Value list of values) instead of the original value. This function is used together with the "match-mapping" MO.

On the compressor side, the match-mapping Matching Operator searches the TV for a match with the header field value and the mapping-sent CDA appends the corresponding index to the Compression Residue to be sent. On the decompressor side, the CDA uses the received index to restore the field value by looking up the list in the TV.

The number of bits sent is the minimal size for coding all the possible indices.

#### 7.5.4. LSB CDA

The LSB action is used together with the "MSB(x)" MO to avoid sending the most significant part of the packet field if that part is already known by the receiving end. The number of bits sent is the original header field length minus the length specified in the MSB(x) MO.

The compressor sends the Least Significant Bits (e.g. LSB of the length field). The decompressor concatenates the x most significant bits of Target Value and the received residue.

If this action needs to be done on a variable length field, the size of the Compression Residue in bytes MUST be sent as described in Section 7.5.

#### 7.5.5. DevIID, AppIID CDA

These functions are used to process respectively the Dev and the App Interface Identifiers (DevIID and AppIID) of the IPv6 addresses. AppIID CDA is less common since current LPWAN technologies frames contain a single address, which is the Dev's address.

The IID value MAY be computed from the Device ID present in the L2 header, or from some other stable identifier. The computation is specific to each LPWAN technology and MAY depend on the Device ID size.

In the downlink direction (Dw), at the compressor, this DevIID CDA may be used to generate the L2 addresses on the LPWAN, based on the packet destination address.

#### 7.5.6. Compute-\*

Some fields are elided during compression and reconstructed during decompression. This is the case for length and checksum, so:

- o compute-length: computes the length assigned to this field. This CDA MAY be used to compute IPv6 length or UDP length.
- o compute-checksum: computes a checksum from the information already received by the SCHC C/D. This field MAY be used to compute UDP checksum.

## 8. Fragmentation

### 8.1. Overview

In LPWAN technologies, the L2 data unit size typically varies from tens to hundreds of bytes. The SCHC F/R (Fragmentation /Reassembly) MAY be used either because after applying SCHC C/D or when SCHC C/D is not possible the entire SCHC Packet still exceeds the L2 data unit.

The SCHC F/R functionality defined in this document has been designed under the assumption that data unit out-of-sequence delivery will not



happen between the entity performing fragmentation and the entity performing reassembly. This assumption allows reducing the complexity and overhead of the SCHC F/R mechanism.

This document also assumes that the L2 data unit size does not vary while a fragmented SCHC Packet is being transmitted.

To adapt the SCHC F/R to the capabilities of LPWAN technologies, it is required to enable optional SCHC Fragment retransmission and to allow for a range of reliability options for sending the SCHC Fragments. This document does not make any decision with regard to which SCHC Fragment delivery reliability mode will be used over a specific LPWAN technology. These details will be defined in other technology-specific documents.

SCHC F/R uses the knowledge of the L2 Word size (see Section 4) to encode some messages. Therefore, SCHC MUST know the L2 Word size. SCHC F/R generates SCHC Fragments and SCHC ACKs that are, for most of them, multiples of L2 Words. The padding overhead is kept to the absolute minimum. See Section 9.

## 8.2. Fragmentation Tools

This subsection describes the different tools that are used to enable the SCHC F/R functionality defined in this document, such as fields in the SCHC F/R header frames (see the related formats in Section 8.4), windows and timers.

- o Rule ID. The Rule ID is present in the SCHC Fragment header and in the SCHC ACK header formats. The Rule ID in a SCHC Fragment header is used to identify that a SCHC Fragment is being carried, which SCHC F/R reliability mode is used and which window size is used. The Rule ID in the SCHC Fragment header also allows interleaving non-fragmented SCHC Packets and SCHC Fragments that carry other SCHC Packets. The Rule ID in a SCHC ACK identifies the message as a SCHC ACK.
- o Fragment Compressed Number (FCN). The FCN is included in all SCHC Fragments. This field can be understood as a truncated, efficient representation of a larger-sized fragment number, and does not carry an absolute SCHC Fragment number. There are two FCN reserved values that are used for controlling the SCHC F/R process, as described next:
  - \* The FCN value with all the bits equal to 1 (All-1) denotes the last SCHC Fragment of a packet. The last window of a packet is called an All-1 window.

- \* The FCN value with all the bits equal to 0 (All-0) denotes the last SCHC Fragment of a window that is not the last one of the packet. Such a window is called an All-0 window.

The rest of the FCN values are assigned in a sequentially decreasing order, which has the purpose to avoid possible ambiguity for the receiver that might arise under certain conditions. In the SCHC Fragments, this field is an unsigned integer, with a size of N bits. In the No-ACK mode, the size is set to 1 bit (N=1), All-0 is used in all SCHC Fragments and All-1 for the last one. For the other reliability modes, it is recommended to use a number of bits (N) equal to or greater than 3. Nevertheless, the appropriate value of N MUST be defined in the corresponding technology-specific profile documents. For windows that are not the last one of a fragmented SCHC Packet, the FCN for the last SCHC Fragment in such windows is an All-0. This indicates that the window is finished and communication proceeds according to the reliability mode in use. The FCN for the last SCHC Fragment in the last window is an All-1, indicating the last SCHC Fragment of the SCHC Packet. It is also important to note that, in the No-ACK mode or when N=1, the last SCHC Fragment of the packet will carry a FCN equal to 1, while all previous SCHC Fragments will carry a FCN to 0. For further details see Section 8.5. The highest FCN in the window, denoted MAX\_WIND\_FCND, MUST be a value equal to or smaller than  $2^N - 2$ . (Example for N=5, MAX\_WIND\_FCND MAY be set to 23, then subsequent FCNs are set sequentially and in decreasing order, and the FCN will wrap from 0 back to 23).

- o Datagram Tag (DTag). The DTag field, if present, is set to the same value for all SCHC Fragments carrying the same SCHC packet, and to different values for different SCHC Packets. Using this field, the sender can interleave fragments from different SCHC Packets, while the receiver can still tell them apart. In the SCHC Fragment formats, the size of the DTag field is T bits, which MAY be set to a value greater than or equal to 0 bits. For each new SCHC Packet processed by the sender, DTag MUST be sequentially increased, from 0 to  $2^T - 1$  wrapping back from  $2^T - 1$  to 0. In the SCHC ACK format, DTag carries the same value as the DTag field in the SCHC Fragments for which this SCHC ACK is intended. When there is no Dtag, there can be only one SCHC Packet in transit. Only after all its fragments have been transmitted can another SCHC Packet be sent. The length of DTag, denoted T, is not specified in this document because it is technology dependant. It will be defined in the corresponding technology-specific documents, based on the number of simultaneous packets that are to be supported.

- o W (window): W is a 1-bit field. This field carries the same value for all SCHC Fragments of a window, and it is complemented for the next window. The initial value for this field is 0. In the SCHC ACK format, this field also has a size of 1 bit. In all SCHC ACKs, the W bit carries the same value as the W bit carried by the SCHC Fragments whose reception is being positively or negatively acknowledged by the SCHC ACK.
- o Message Integrity Check (MIC). This field is computed by the sender over the complete SCHC Packet and before SCHC fragmentation. The MIC allows the receiver to check errors in the reassembled packet, while it also enables compressing the UDP checksum by use of SCHC compression. The CRC32 as 0xEDB88320 (i.e. the reverse representation of the polynomial used e.g. in the Ethernet standard [RFC3385]) is recommended as the default algorithm for computing the MIC. Nevertheless, other algorithms MAY be required and are defined in the technology-specific documents as well as the length in bits of the MIC used.
- o C (MIC checked): C is a 1-bit field. This field is used in the SCHC ACK packets to report the outcome of the MIC check, i.e. whether the reassembled packet was correctly received or not. A value of 1 represents a positive MIC check at the receiver side (i.e. the MIC computed by the receiver matches the received MIC).
- o Retransmission Timer. A SCHC Fragment sender uses it after the transmission of a window to detect a transmission error of the SCHC ACK corresponding to this window. Depending on the reliability mode, it will lead to a request a SCHC ACK retransmission (in ACK-Always mode) or it will trigger the transmission of the next window (in ACK-on-Error mode). The duration of this timer is not defined in this document and MUST be defined in the corresponding technology-specific documents.
- o Inactivity Timer. A SCHC Fragment receiver uses it to take action when there is a problem in the transmission of SCHC fragments. Such a problem could be detected by the receiver not getting a single SCHC Fragment during a given period of time. When this happens, an Abort message will be sent (see related text later in this section). Initially, and each time a SCHC Fragment is received, the timer is reinitialized. The duration of this timer is not defined in this document and MUST be defined in the corresponding technology-specific document.
- o Attempts. This counter counts the requests for a missing SCHC ACK. When it reaches the value MAX\_ACK\_REQUESTS, the sender assumes there are recurrent SCHC Fragment transmission errors and determines that an Abort is needed. The default value

MAX\_ACK\_REQUESTS is not stated in this document, and it is expected to be defined in the corresponding technology-specific document. The Attempts counter is defined per window. It is initialized each time a new window is used.

- o Bitmap. The Bitmap is a sequence of bits carried in a SCHC ACK. Each bit in the Bitmap corresponds to a SCHC fragment of the current window, and provides feedback on whether the SCHC Fragment has been received or not. The right-most position on the Bitmap reports if the All-0 or All-1 fragment has been received or not. Feedback on the SCHC fragment with the highest FCN value is provided by the bit in the left-most position of the Bitmap. In the Bitmap, a bit set to 1 indicates that the SCHC Fragment of FCN corresponding to that bit position has been correctly sent and received. The text above describes the internal representation of the Bitmap. When inserted in the SCHC ACK for transmission from the receiver to the sender, the Bitmap is shortened for energy/bandwidth optimisation, see more details in Section 8.4.3.1.
- o Abort. On expiration of the Inactivity timer, or when Attempts reaches MAX\_ACK\_REQUESTS or upon occurrence of some other error, the sender or the receiver may use the Abort. When the receiver needs to abort the on-going fragmented SCHC Packet transmission, it sends the Receiver-Abort format. When the sender needs to abort the transmission, it sends the Sender-Abort format. None of the Aborts are acknowledged.

### 8.3. Reliability modes

This specification defines three reliability modes: No-ACK, ACK-Always, and ACK-on-Error. ACK-Always and ACK-on-Error operate on windows of SCHC Fragments. A window of SCHC Fragments is a subset of the full set of SCHC Fragments needed to carry a SCHC Packet.

- o No-ACK. No-ACK is the simplest SCHC Fragment reliability mode. The receiver does not generate overhead in the form of acknowledgements (ACKs). However, this mode does not enhance reliability beyond that offered by the underlying LPWAN technology. In the No-ACK mode, the receiver MUST NOT issue SCHC ACKs. See further details in Section 8.5.1.
- o ACK-Always. The ACK-Always mode provides flow control using a windowing scheme. This mode is also able to handle long bursts of lost SCHC Fragments since detection of such events can be done before the end of the SCHC Packet transmission as long as the window size is short enough. However, such benefit comes at the expense of SCHC ACK use. In ACK-Always, the receiver sends a SCHC ACK after a window of SCHC Fragments has been received. The SCHC

ACK is used to inform the sender which SCHC Fragments in the current window have been well received. Upon a SCHC ACK reception, the sender retransmits the lost SCHC Fragments. When a SCHC ACK is lost and the sender has not received it by the expiration of the Retransmission Timer, the sender uses a SCHC ACK request by sending the All-0 empty SCHC Fragment when it is not the last window and the All-1 empty Fragment when it is the last window. The maximum number of SCHC ACK requests is `MAX_ACK_REQUESTS`. If `MAX_ACK_REQUESTS` is reached, the transmission needs to be aborted. See further details in Section 8.5.2.

- o ACK-on-Error. The ACK-on-Error mode is suitable for links offering relatively low L2 data unit loss probability. In this mode, the SCHC Fragment receiver reduces the number of SCHC ACKs transmitted, which MAY be especially beneficial in asymmetric scenarios. The receiver transmits a SCHC ACK only after the complete window transmission and if at least one SCHC Fragment of this window has been lost. An exception to this behavior is in the last window, where the receiver MUST transmit a SCHC ACK, including the C bit set based on the MIC checked result, even if all the SCHC Fragments of the last window have been correctly received. The SCHC ACK gives the state of all the SCHC Fragments of the current window (received or lost). Upon a SCHC ACK reception, the sender retransmits any lost SCHC Fragments based on the SCHC ACK. If a SCHC ACK is not transmitted back by the receiver at the end of a window, the sender assumes that all SCHC Fragments have been correctly received. When a SCHC ACK is lost, the sender assumes that all SCHC Fragments covered by the lost SCHC ACK have been successfully delivered, so the sender continues transmitting the next window of SCHC Fragments. If the next SCHC Fragments received belong to the next window and it is still expecting fragments from the previous window, the receiver will abort the on-going fragmented packet transmission. See further details in Section 8.5.3.

The same reliability mode MUST be used for all SCHC Fragments of a SCHC Packet. The decision on which reliability mode will be used and whether the same reliability mode applies to all SCHC Packets is an implementation problem and is out of the scope of this document.

Note that the reliability mode choice is not necessarily tied to a particular characteristic of the underlying L2 LPWAN technology, e.g. the No-ACK mode MAY be used on top of an L2 LPWAN technology with symmetric characteristics for uplink and downlink. This document does not make any decision as to which SCHC Fragment reliability modes are relevant for a specific LPWAN technology.

Examples of the different reliability modes described are provided in Appendix B.

8.4. Fragmentation Formats

This section defines the SCHC Fragment format, including the All-0 and All-1 formats and their "empty" variations, the SCHC ACK format and the Abort formats.

A SCHC Fragment conforms to the general format shown in Figure 11. It comprises a SCHC Fragment Header and a SCHC Fragment Payload. In addition, the last SCHC Fragment carries as many padding bits as needed to fill up an L2 Word. The SCHC Fragment Payload carries a subset of the SCHC Packet. The SCHC Fragment is the data unit passed on to the L2 for transmission.

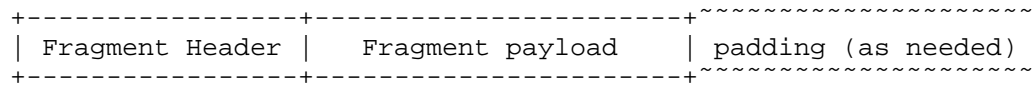


Figure 11: SCHC Fragment general format. Presence of a padding field is optional

8.4.1. Fragments that are not the last one

In ACK-Always or ACK-on-Error, SCHC Fragments except the last one SHALL conform to the detailed format defined in Figure 12.




Figure 12: Fragment Detailed Format for Fragments except the Last One, ACK-Always and ACK-on-Error

In the No-ACK mode, SCHC Fragments except the last one SHALL conform to the detailed format defined in Figure 13.

```

|---- Fragment Header ----|
      |-- T --|-- N --|
+-- ... ---+ ... +- ... -+-----+
| Rule ID | DTag | FCN | Fragment payload |
+-- ... ---+ ... +- ... -+-----+

```

Figure 13: Fragment Detailed Format for Fragments except the Last One, No-ACK mode

The total size of the fragment header is not necessarily a multiple of the L2 Word size. To build the fragment payload, SCHC F/R MUST take from the SCHC Packet a number of bits that makes the SCHC Fragment an exact multiple of L2 Words. As a consequence, no padding bit is used for these fragments.

#### 8.4.1.1. All-0 fragment

The All-0 format is used for sending the last SCHC Fragment of a window that is not the last window of the SCHC Packet.

```

|----- Fragment Header -----|
      |-- T --|1|-- N --|
+-- ... ---+ ... +---+ ... -+-----+
| Rule ID | DTag | W | 0..0 | Fragment payload |
+-- ... ---+ ... +---+ ... -+-----+

```

Figure 14: All-0 fragment detailed format

This is simply an instance of the format described in Figure 12. An All-0 fragment payload MUST be at least the size of an L2 Word. The rationale is that the All-0 empty fragment (see Section 8.4.1.2) needs to be distinguishable from the All-0 regular fragment, even in the presence of padding.

#### 8.4.1.2. All-0 empty fragment

The All-0 empty fragment is an exception to the All-0 fragment described above. It is used by a sender to request the retransmission of a SCHC ACK by the receiver. It is only used in ACK-Always mode.




Figure 15: All-0 empty fragment detailed format

The size of the All-0 fragment header is generally not a multiple of the L2 Word size. Therefore, an All-0 empty fragment generally needs padding bits. The padding bits are always less than an L2 Word.

Since an All-0 payload MUST be at least the size of an L2 Word, a receiver can distinguish an All-0 empty fragment from a regular All-0 fragment, even in the presence of padding.

8.4.2. All-1 fragment

In the No-ACK mode, the last SCHC Fragment of a SCHC Packet SHALL contain a SCHC Fragment header that conforms to the detailed format shown in Figure 16.




Figure 16: All-1 Fragment Detailed Format for the Last Fragment, No-ACK mode

In ACK-Always or ACK-on-Error mode, the last fragment of a SCHC Packet SHALL contain a SCHC Fragment header that conforms to the detailed format shown in Figure 17.




Figure 17: All-1 Fragment Detailed Format for the Last Fragment, ACK-Always or ACK-on-Error



The total size of the All-1 SCHC Fragment header is generally not a multiple of the L2 Word size. The All-1 fragment being the last one of the SCHC Packet, SCHC F/R cannot freely choose the payload size to align the fragment to an L2 Word. Therefore, padding bits are generally appended to the All-1 fragment to make it a multiple of L2 Words in size.

The MIC MUST be computed on the payload and the padding bits. The rationale is that the SCHC Reassembler needs to check the correctness of the reassembled SCHC packet but has no way of knowing where the payload ends. Indeed, the latter requires decompressing the SCHC Packet.

An All-1 fragment payload MUST be at least the size of an L2 Word. The rationale is that the All-1 empty fragment (see Section 8.4.2.1) needs to be distinguishable from the All-1 fragment, even in the presence of padding. This may entail saving an L2 Word from the previous fragment payload to make the payload of this All-1 fragment big enough.

The values for N, T and the length of MIC are not specified in this document, and SHOULD be determined in other documents (e.g. technology-specific profile documents).

The length of the MIC MUST be at least an L2 Word size. The rationale is to be able to distinguish a Sender-Abort (see Section 8.4.4) from an All-1 Fragment, even in the presence of padding.

#### 8.4.2.1. All-1 empty fragment

The All-1 empty fragment format is an All-1 fragment format without a payload (see Figure 18). It is used by a fragment sender, in either ACK-Always or ACK-on-Error, to request a retransmission of the SCHC ACK for the All-1 window.

The size of the All-1 empty fragment header is generally not a multiple of the L2 Word size. Therefore, an All-1 empty fragment generally needs padding bits. The padding bits are always less than an L2 Word.

Since an All-1 payload MUST be at least the size of an L2 Word, a receiver can distinguish an All-1 empty fragment from a regular All-1 fragment, even in the presence of padding.




Figure 18: All-1 for Retries format, also called All-1 empty

8.4.3. SCHC ACK format

The format of a SCHC ACK that acknowledges a window that is not the last one (denoted as All-0 window) is shown in Figure 19.




Figure 19: ACK format for All-0 windows

To acknowledge the last window of a packet (denoted as All-1 window), a C bit (i.e. MIC checked) following the W bit is set to 1 to indicate that the MIC check computed by the receiver matches the MIC present in the All-1 fragment. If the MIC check fails, the C bit is set to 0 and the Bitmap for the All-1 window follows.




Figure 20: Format of a SCHC ACK for All-1 windows

The Rule ID and Dtag values in the SCHC ACK messages MUST be identical to the ones used in the SCHC Fragments that are being acknowledged. This allows matching the SCHC ACK and the corresponding SCHC Fragments.

The Bitmap carries information on the reception of each fragment of the window as described in Section 8.2.

See Appendix D for a discussion on the size of the Bitmaps.

In order to reduce the SCK ACK size, the Bitmap that is actually transmitted is shortened ("encoded") as explained in Section 8.4.3.1.

### 8.4.3.1. Bitmap Encoding

The SCHC ACK that is transmitted is truncated by applying the following algorithm: the longest contiguous sequence of bits that starts at an L2 Word boundary of the SCHC ACK, where the bits of that sequence are all set to 1, are all part of the Bitmap and finish exactly at the end of the Bitmap, if one such sequence exists, **MUST NOT** be transmitted. Because the SCHC Fragment sender knows the actual Bitmap size, it can reconstruct the original Bitmap from the shortened bitmap.

When shortening effectively takes place, the SCHC ACK is a multiple of L2 Words, and padding **MUST NOT** be appended. When shortening does not happen, padding bits **MUST** be appended as needed to fill up the last L2 Word.

Figure 21 shows an example where L2 Words are actually bytes and where the original Bitmap contains 17 bits, the last 15 of which are all set to 1.




Figure 21: A non-encoded Bitmap

Figure 22 shows that the last 14 bits are not sent.




Figure 22: Optimized Bitmap format

Figure 23 shows an example of a SCHC ACK with FCN ranging from 6 down to 0, where the Bitmap indicates that the second and the fifth SCHC Fragments have not been correctly received.

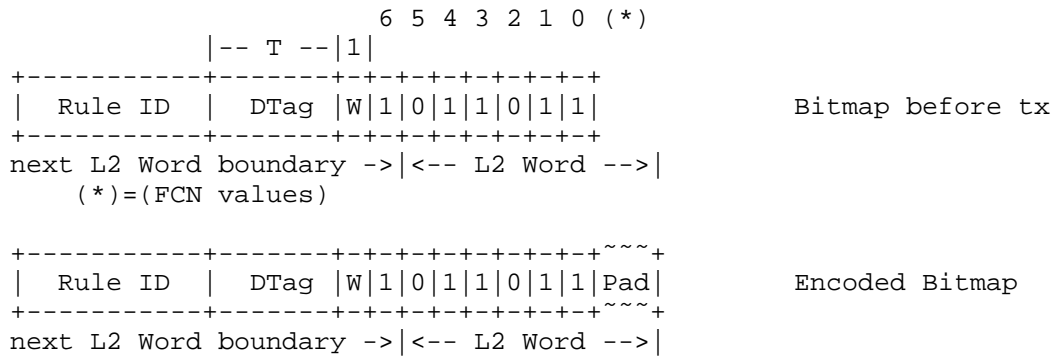


Figure 23: Example of a Bitmap before transmission, and the transmitted one, for a window that is not the last one

Figure 24 shows an example of a SCHC ACK with FCN ranging from 6 down to 0, where MIC check has failed but the Bitmap indicates that there is no missing SCHC Fragment.

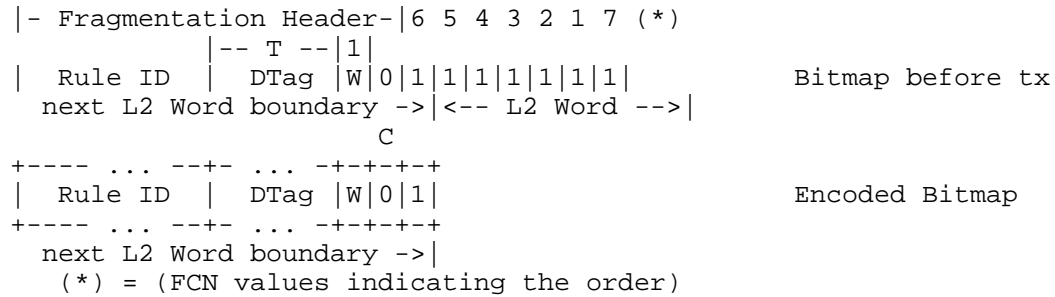


Figure 24: Example of the Bitmap in ACK-Always or ACK-on-Error for the last window

#### 8.4.4. Abort formats

When a SCHC Fragment sender needs to abort the on-going fragmented SCHC Packet transmission, it sends a Sender-Abort. The Sender-Abort format (see Figure 25) is a variation of the All-1 fragment, with neither a MIC nor a payload. All-1 fragments contain at least a MIC. The absence of the MIC indicates a Sender-Abort.

```
|--- Sender-Abort Header ---|
+--- ... ----+ ... -++-...-+~
| Rule ID | DTag |W| FCN | padding (as needed)
+--- ... ----+ ... -++-...-+~
```

Figure 25: Sender-Abort format. All FCN field bits in this format are set to 1

The size of the Sender-Abort header is generally not a multiple of the L2 Word size. Therefore, a Sender-Abort generally needs padding bits.

Since an All-1 fragment MIC MUST be at least the size of an L2 Word, a receiver can distinguish a Sender-Abort from an All-1 fragment, even in the presence of padding.

When a SCHC Fragment receiver needs to abort the on-going fragmented SCHC Packet transmission, it transmits a Receiver-Abort. The Receiver-Abort format is a variation on the SCHC ACK format, creating an exception in the encoded Bitmap algorithm. As shown in Figure 26, a Receiver-Abort is coded as a SCHC ACK message with a shortened Bitmap set to 1 up to the first L2 Word boundary, followed by an extra L2 Word full of 1's. Such a message never occurs in a regular acknowledgement and is detected as a Receiver-Abort.

The Rule ID and Dtag values in the Receive-Abort message MUST be identical to the ones used in the fragments of the SCHC Packet the transmission of which is being aborted.

A Receiver-Abort is aligned to L2 Words, by design. Therefore, padding MUST NOT be appended.

```
|- Receiver-Abort Header -|
+---- ... ----+ ... -++-...-+~
| Rule ID | DTag |W| 1..1| 1..1 |
+---- ... ----+ ... -++-...-+~
      next L2 Word boundary ->|<-- L2 Word -->|
```

Figure 26: Receiver-Abort format

Neither the Sender-Abort nor the Receiver-Abort messages are ever acknowledged or retransmitted.

Use cases for the Sender-Abort and Receiver-Abort messages are explained in Section 8.5 or Appendix C.

### 8.5. Baseline mechanism

If after applying SCHC header compression (or when SCHC header compression is not possible) the SCHC Packet does not fit within the payload of a single L2 data unit, the SCHC Packet SHALL be broken into SCHC Fragments and the fragments SHALL be sent to the fragment receiver. The fragment receiver needs to identify all the SCHC Fragments that belong to a given SCHC Packet. To this end, the receiver SHALL use:

- o The sender's L2 source address (if present),
- o The destination's L2 address (if present),
- o Rule ID,
- o DTag (if present).

Then, the fragment receiver MAY determine the SCHC Fragment reliability mode that is used for this SCHC Fragment based on the Rule ID in that fragment.

After a SCHC Fragment reception, the receiver starts constructing the SCHC Packet. It uses the FCN and the arrival order of each SCHC Fragment to determine the location of the individual fragments within the SCHC Packet. For example, the receiver MAY place the fragment payload within a payload reassembly buffer at the location determined from the FCN, the arrival order of the SCHC Fragments, and the fragment payload sizes. In ACK-on-Error or ACK-Always, the fragment receiver also uses the W bit in the received SCHC Fragments. Note that the size of the original, unfragmented packet cannot be determined from fragmentation headers.

Fragmentation functionality uses the FCN value to transmit the SCHC Fragments. It has a length of N bits where the All-1 and All-0 FCN values are used to control the fragmentation transmission. The rest of the FCN numbers MUST be assigned sequentially in a decreasing order, the first FCN of a window is RECOMMENDED to be MAX\_WIND\_FCN, i.e. the highest possible FCN value depending on the FCN number of bits.

In all modes, the last SCHC Fragment of a packet MUST contain a MIC which is used to check if there are errors or missing SCHC Fragments and MUST use the corresponding All-1 fragment format. Note that a SCHC Fragment with an All-0 format is considered the last SCHC Fragment of the current window.

If the receiver receives the last fragment of a SCHC Packet (All-1), it checks for the integrity of the reassembled SCHC Packet, based on the MIC received. In No-ACK, if the integrity check indicates that the reassembled SCHC Packet does not match the original SCHC Packet (prior to fragmentation), the reassembled SCHC Packet MUST be discarded. In ACK-on-Error or ACK-Always, a MIC check is also performed by the fragment receiver after reception of each subsequent SCHC Fragment retransmitted after the first MIC check.

Notice that the SCHC ACK for the All-1 window carries one more bit (the C bit) compared to the SCHC ACKs for the previous windows. See Appendix D for a discussion on various options to deal with this "bump" in the SCHC ACK.

There are three reliability modes: No-ACK, ACK-Always and ACK-on-Error. In ACK-Always and ACK-on-Error, a jumping window protocol uses two windows alternatively, identified as 0 and 1. A SCHC Fragment with all FCN bits set to 0 (i.e. an All-0 fragment) indicates that the window is over (i.e. the SCHC Fragment is the last one of the window) and allows to switch from one window to the next one. The All-1 FCN in a SCHC Fragment indicates that it is the last fragment of the packet being transmitted and therefore there will not be another window for this packet.

#### 8.5.1. No-ACK

In the No-ACK mode, there is no feedback communication from the fragment receiver. The sender will send all the SCHC fragments of a packet without any possibility of knowing if errors or losses have occurred. As, in this mode, there is no need to identify specific SCHC Fragments, a one-bit FCN MAY be used. Consequently, the FCN All-0 value is used in all SCHC fragments except the last one, which carries an All-1 FCN and the MIC. The receiver will wait for SCHC Fragments and will set the Inactivity timer. The receiver will use the MIC contained in the last SCHC Fragment to check for errors. When the Inactivity Timer expires or if the MIC check indicates that the reassembled packet does not match the original one, the receiver will release all resources allocated to reassembling this packet. The initial value of the Inactivity Timer will be determined based on the characteristics of the underlying LPWAN technology and will be defined in other documents (e.g. technology-specific profile documents).

#### 8.5.2. ACK-Always

In ACK-Always, the sender transmits SCHC Fragments by using the two-jumping-windows procedure. A delay between each SCHC fragment can be added to respect local regulations or other constraints imposed by

the applications. Each time a SCHC fragment is sent, the FCN is decreased by one. When the FCN reaches value 0, if there are more SCHC Fragments remaining to be sent, the sender transmits the last SCHC Fragment of this window using the All-0 fragment format. It then starts the Retransmission Timer and waits for a SCHC ACK. Otherwise, if FCN reaches 0 and the sender transmits the last SCHC Fragment of the SCHC Packet, the sender uses the All-1 fragment format, which includes a MIC. The sender sets the Retransmission Timer and waits for the SCHC ACK to know if transmission errors have occurred.

The Retransmission Timer is dimensioned based on the LPWAN technology in use. When the Retransmission Timer expires, the sender sends an All-0 empty (resp. All-1 empty) fragment to request again the SCHC ACK for the window that ended with the All-0 (resp. All-1) fragment just sent. The window number is not changed.

After receiving an All-0 or All-1 fragment, the receiver sends a SCHC ACK with an encoded Bitmap reporting whether any SCHC fragments have been lost or not. When the sender receives a SCHC ACK, it checks the W bit carried by the SCHC ACK. Any SCHC ACK carrying an unexpected W bit value is discarded. If the W bit value of the received SCHC ACK is correct, the sender analyzes the rest of the SCHC ACK message, such as the encoded Bitmap and the MIC. If all the SCHC Fragments sent for this window have been well received, and if at least one more SCHC Fragment needs to be sent, the sender advances its sending window to the next window value and sends the next SCHC Fragments. If no more SCHC Fragments have to be sent, then the fragmented SCHC Packet transmission is finished.

However, if one or more SCHC Fragments have not been received as per the SCHC ACK (i.e. the corresponding bits are not set in the encoded Bitmap) then the sender resends the missing SCHC Fragments. When all missing SCHC Fragments have been retransmitted, the sender starts the Retransmission Timer, even if an All-0 or an All-1 has not been sent as part of this retransmission and waits for a SCHC ACK. Upon receipt of the SCHC ACK, if one or more SCHC Fragments have not yet been received, the counter Attempts is increased and the sender resends the missing SCHC Fragments again. When Attempts reaches MAX\_ACK\_REQUESTS, the sender aborts the on-going fragmented SCHC Packet transmission by sending a Sender-Abort message and releases any resources for transmission of the packet. The sender also aborts an on-going fragmented SCHC Packet transmission when a failed MIC check is reported by the receiver or when a SCHC Fragment that has not been sent is reported in the encoded Bitmap.

On the other hand, at the beginning, the receiver side expects to receive window 0. Any SCHC Fragment received but not belonging to



the current window is discarded. All SCHC Fragments belonging to the correct window are accepted, and the actual SCHC Fragment number managed by the receiver is computed based on the FCN value. The receiver prepares the encoded Bitmap to report the correctly received and the missing SCHC Fragments for the current window. After each SCHC Fragment is received, the receiver initializes the Inactivity Timer. When the Inactivity Timer expires, the transmission is aborted by the receiver sending a Receiver-Abort message.

When an All-0 fragment is received, it indicates that all the SCHC Fragments have been sent in the current window. Since the sender is not obliged to always send a full window, some SCHC Fragment number not set in the receiver memory may not correspond to losses. The receiver sends the corresponding SCHC ACK, the Inactivity Timer is set and the transmission of the next window by the sender can start.

If an All-0 fragment has been received and all SCHC Fragments of the current window have also been received, the receiver then expects a new Window and waits for the next SCHC Fragment. Upon receipt of a SCHC Fragment, if the window value has not changed, the received SCHC Fragments are part of a retransmission. A receiver that has already received a SCHC Fragment SHOULD discard it, otherwise, it updates the Bitmap. If all the bits of the Bitmap are set to one, the receiver MUST send a SCHC ACK without waiting for an All-0 fragment and the Inactivity Timer is initialized.

On the other hand, if the window value of the next received SCHC Fragment is set to the next expected window value, this means that the sender has received a correct encoded Bitmap reporting that all SCHC Fragments have been received. The receiver then updates the value of the next expected window.

When an All-1 fragment is received, it indicates that the last SCHC Fragment of the packet has been sent. Since the last window is not always full, the MIC will be used by the receiver to detect if all SCHC Fragments of the packet have been received. A correct MIC indicates the end of the transmission but the receiver MUST stay alive for an Inactivity Timer period to answer to any empty All-1 fragments the sender MAY send if SCHC ACKs sent by the receiver are lost. If the MIC is incorrect, some SCHC Fragments have been lost. The receiver sends the SCHC ACK regardless of successful fragmented SCHC Packet reception or not, the Inactivity Timer is set. In case of an incorrect MIC, the receiver waits for SCHC Fragments belonging to the same window. After `MAX_ACK_REQUESTS`, the receiver will abort the on-going fragmented SCHC Packet transmission by transmitting a the Receiver-Abort format. The receiver also aborts upon Inactivity Timer expiration by sending a Receiver-Abort message.

If the sender receives a SCK ACK with a Bitmap containing a bit set for a SCHC Fragment that it has not sent during the transmission phase of this window, it MUST abort the whole fragmentation and transmission of this SCHC Packet.

### 8.5.3. ACK-on-Error

The senders behavior for ACK-on-Error and ACK-Always are similar. The main difference is that in ACK-on-Error the SCHC ACK with the encoded Bitmap is not sent at the end of each window but only when at least one SCHC Fragment of the current window has been lost. Except for the last window where a SCHC ACK MUST be sent to finish the transmission.

In ACK-on-Error, the Retransmission Timer expiration is considered as a positive acknowledgement for all windows but the last one. This timer is set after sending an All-0 or an All-1 fragment. For an All-0 fragment, on timer expiration, the sender resumes operation and sends the SCHC Fragments of the next window.

If the sender receives a SCHC ACK, it checks the window value. SCHC ACKs with an unexpected window number are discarded. If the window number in the received SCHC ACK is correct, the sender verifies if the receiver has received all SCHC fragments of the current window. When at least one SCHC Fragment has been lost, the counter Attempts is increased by one and the sender resends the missing SCHC Fragments again. When Attempts reaches MAX\_ACK\_REQUESTS, the sender sends a Sender-Abort message and releases all resources for the on-going fragmented SCHC Packet transmission. When the retransmission of the missing SCHC Fragments is finished, the sender starts listening for a SCHC ACK (even if an All-0 or an All-1 has not been sent during the retransmission) and initializes the Retransmission Timer.

After sending an All-1 fragment, the sender listens for a SCHC ACK, initializes Attempts, and starts the Retransmission Timer. If the Retransmission Timer expires, Attempts is increased by one and an empty All-1 fragment is sent to request the SCHC ACK for the last window. If Attempts reaches MAX\_ACK\_REQUESTS, the sender aborts the on-going fragmented SCHC Packet transmission by transmitting the Sender-Abort fragment.

At the end of any window, if the sender receives a SCK ACK with a Bitmap containing a bit set for a SCHC Fragment that it has not sent during the transmission phase of that window, it MUST abort the whole fragmentation and transmission of this SCHC Packet.

Unlike the sender, the receiver for ACK-on-Error has a larger amount of differences compared with ACK-Always. First, a SCHC ACK is not

sent unless there is a lost SCHC Fragment or an unexpected behavior. With the exception of the last window, where a SCHC ACK is always sent regardless of SCHC Fragment losses or not. The receiver starts by expecting SCHC Fragments from window 0 and maintains the information regarding which SCHC Fragments it receives. After receiving a SCHC Fragment, the Inactivity Timer is set. If no further SCHC Fragment are received and the Inactivity Timer expires, the SCHC Fragment receiver aborts the on-going fragmented SCHC Packet transmission by transmitting the Receiver-Abort data unit.

Any SCHC Fragment not belonging to the current window is discarded. The actual SCHC Fragment number is computed based on the FCN value. When an All-0 fragment is received and all SCHC Fragments have been received, the receiver updates the expected window value and expects a new window and waits for the next SCHC Fragment. If the window value of the next SCHC Fragment has not changed, the received SCHC Fragment is a retransmission. A receiver that has already received a Fragment discard it. If all SCHC Fragments of a window (that is not the last one) have been received, the receiver does not send a SCHC ACK. While the receiver waits for the next window and if the window value is set to the next value, and if an All-1 fragment with the next value window arrived the receiver knows that the last SCHC Fragment of the packet has been sent. Since the last window is not always full, the MIC will be used to detect if all SCHC Fragments of the window have been received. A correct MIC check indicates the end of the fragmented SCHC Packet transmission. An ACK is sent by the SCHC Fragment receiver. In case of an incorrect MIC, the receiver waits for SCHC Fragments belonging to the same window or the expiration of the Inactivity Timer. The latter will lead the receiver to abort the on-going SCHC fragmented packet transmission by transmitting the Receiver-Abort message.

If, after receiving an All-0 fragment the receiver missed some SCHC Fragments, the receiver uses a SCHC ACK with the encoded Bitmap to ask the retransmission of the missing fragments and expect to receive SCHC Fragments with the actual window. While waiting the retransmission an All-0 empty fragment is received, the receiver sends again the SCHC ACK with the encoded Bitmap, if the SCHC Fragments received belongs to another window or an All-1 fragment is received, the transmission is aborted by sending a Receiver-Abort fragment. Once it has received all the missing fragments it waits for the next window fragments.

#### 8.6. Supporting multiple window sizes

For ACK-Always or ACK-on-Error, implementers MAY opt to support a single window size or multiple window sizes. The latter, when feasible, may provide performance optimizations. For example, a

large window size SHOULD be used for packets that need to be carried by a large number of SCHC Fragments. However, when the number of SCHC Fragments required to carry a packet is low, a smaller window size, and thus a shorter Bitmap, MAY be sufficient to provide feedback on all SCHC Fragments. If multiple window sizes are supported, the Rule ID MAY be used to signal the window size in use for a specific packet transmission.

Note that the same window size MUST be used for the transmission of all SCHC Fragments that belong to the same SCHC Packet.

#### 8.7. Downlink SCHC Fragment transmission

In some LPWAN technologies, as part of energy-saving techniques, downlink transmission is only possible immediately after an uplink transmission. In order to avoid potentially high delay in the downlink transmission of a fragmented SCHC Packet, the SCHC Fragment receiver MAY perform an uplink transmission as soon as possible after reception of a SCHC Fragment that is not the last one. Such uplink transmission MAY be triggered by the L2 (e.g. an L2 ACK sent in response to a SCHC Fragment encapsulated in a L2 frame that requires an L2 ACK) or it MAY be triggered from an upper layer.

For downlink transmission of a fragmented SCHC Packet in ACK-Always mode, the SCHC Fragment receiver MAY support timer-based SCHC ACK retransmission. In this mechanism, the SCHC Fragment receiver initializes and starts a timer (the Inactivity Timer is used) after the transmission of a SCHC ACK, except when the SCHC ACK is sent in response to the last SCHC Fragment of a packet (All-1 fragment). In the latter case, the SCHC Fragment receiver does not start a timer after transmission of the SCHC ACK.

If, after transmission of a SCHC ACK that is not an All-1 fragment, and before expiration of the corresponding Inactivity timer, the SCHC Fragment receiver receives a SCHC Fragment that belongs to the current window (e.g. a missing SCHC Fragment from the current window) or to the next window, the Inactivity timer for the SCHC ACK is stopped. However, if the Inactivity timer expires, the SCHC ACK is resent and the Inactivity timer is reinitialized and restarted.

The default initial value for the Inactivity timer, as well as the maximum number of retries for a specific SCHC ACK, denoted `MAX_ACK_RETRIES`, are not defined in this document, and need to be defined in other documents (e.g. technology-specific profiles). The initial value of the Inactivity timer is expected to be greater than that of the Retransmission timer, in order to make sure that a (buffered) SCHC Fragment to be retransmitted can find an opportunity for that transmission.

When the SCHC Fragment sender transmits the All-1 fragment, it starts its Retransmission Timer with a large timeout value (e.g. several times that of the initial Inactivity timer). If a SCHC ACK is received before expiration of this timer, the SCHC Fragment sender retransmits any lost SCHC Fragments reported by the SCHC ACK, or if the SCHC ACK confirms successful reception of all SCHC Fragments of the last window, the transmission of the fragmented SCHC Packet is considered complete. If the timer expires, and no SCHC ACK has been received since the start of the timer, the SCHC Fragment sender assumes that the All-1 fragment has been successfully received (and possibly, the last SCHC ACK has been lost: this mechanism assumes that the retransmission timer for the All-1 fragment is long enough to allow several SCHC ACK retries if the All-1 fragment has not been received by the SCHC Fragment receiver, and it also assumes that it is unlikely that several ACKs become all lost).

#### 9. Padding management

SCHC C/D and SCHC F/R operate on bits, not bytes. SCHC itself does not have any alignment prerequisite. If the Layer 2 below SCHC constrains the L2 Data Unit to align to some boundary, called L2 Words (for example, bytes), SCHC will meet that constraint and produce messages with the correct alignment. This may entail adding extra bits (called padding bits).

When padding occurs, the number of appended bits is strictly less than the L2 Word size.

Padding happens at most once for each Packet going through the full SCHC chain, i.e. Compression and (optionally) SCHC Fragmentation (see Figure 2). If a SCHC Packet is sent unfragmented (see Figure 27), it is padded as needed. If a SCHC Packet is fragmented, only the last fragment is padded as needed.

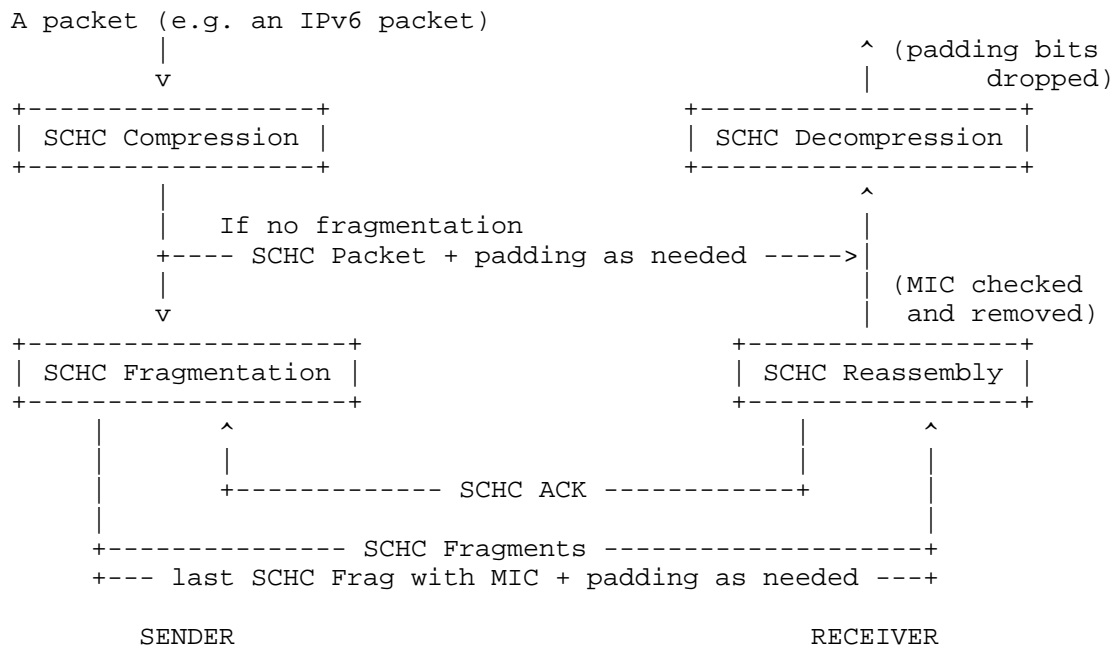


Figure 27: SCHC operations, including padding as needed

Each technology-specific document MUST specify the size of the L2 Word. The L2 Word might actually be a single bit, in which case at most zero bits of padding will be appended to any message, i.e. no padding will take place at all.

10. SCHC Compression for IPv6 and UDP headers

This section lists the different IPv6 and UDP header fields and how they can be compressed.

10.1. IPv6 version field

This field always holds the same value. Therefore, in the Rule, TV is set to 6, MO to "equal" and CDA to "not-sent".

10.2. IPv6 Traffic class field

If the DiffServ field does not vary and is known by both sides, the Field Descriptor in the Rule SHOULD contain a TV with this well-known value, an "equal" MO and a "not-sent" CDA.

Otherwise, two possibilities can be considered depending on the variability of the value:

- o One possibility is to not compress the field and send the original value. In the Rule, TV is not set to any particular value, MO is set to "ignore" and CDA is set to "value-sent".
- o If some upper bits in the field are constant and known, a better option is to only send the LSBs. In the Rule, TV is set to a value with the stable known upper part, MO is set to MSB(x) and CDA to LSB(y).

### 10.3. Flow label field

If the Flow Label field does not vary and is known by both sides, the Field Descriptor in the Rule SHOULD contain a TV with this well-known value, an "equal" MO and a "not-sent" CDA.

Otherwise, two possibilities can be considered:

- o One possibility is to not compress the field and send the original value. In the Rule, TV is not set to any particular value, MO is set to "ignore" and CDA is set to "value-sent".
- o If some upper bits in the field are constant and known, a better option is to only send the LSBs. In the Rule, TV is set to a value with the stable known upper part, MO is set to MSB(x) and CDA to LSB(y).

### 10.4. Payload Length field

This field can be elided for the transmission on the LPWAN network. The SCHC C/D recomputes the original payload length value. In the Field Descriptor, TV is not set, MO is set to "ignore" and CDA is "compute-IPv6-length".

If the payload length needs to be sent and does not need to be coded in 16 bits, the TV can be set to 0x0000, the MO set to MSB(16-s) where 's' is the number of bits to code the maximum length, and CDA is set to LSB(s).

### 10.5. Next Header field

If the Next Header field does not vary and is known by both sides, the Field Descriptor in the Rule SHOULD contain a TV with this Next Header value, the MO SHOULD be "equal" and the CDA SHOULD be "not-sent".

Otherwise, TV is not set in the Field Descriptor, MO is set to "ignore" and CDA is set to "value-sent". Alternatively, a matching-list MAY also be used.

#### 10.6. Hop Limit field

The field behavior for this field is different for Uplink and Downlink. In Uplink, since there is no IP forwarding between the Dev and the SCHC C/D, the value is relatively constant. On the other hand, the Downlink value depends of Internet routing and MAY change more frequently. One neat way of processing this field is to use the Direction Indicator (DI) to distinguish both directions:

- o in the Uplink, elide the field: the TV in the Field Descriptor is set to the known constant value, the MO is set to "equal" and the CDA is set to "not-sent".
- o in the Downlink, send the value: TV is not set, MO is set to "ignore" and CDA is set to "value-sent".

#### 10.7. IPv6 addresses fields

As in 6LOWPAN [RFC4944], IPv6 addresses are split into two 64-bit long fields; one for the prefix and one for the Interface Identifier (IID). These fields SHOULD be compressed. To allow for a single Rule being used for both directions, these values are identified by their role (DEV or APP) and not by their position in the frame (source or destination).

##### 10.7.1. IPv6 source and destination prefixes

Both ends MUST be synchronized with the appropriate prefixes. For a specific flow, the source and destination prefixes can be unique and stored in the context. It can be either a link-local prefix or a global prefix. In that case, the TV for the source and destination prefixes contain the values, the MO is set to "equal" and the CDA is set to "not-sent".

If the Rule is intended to compress packets with different prefix values, match-mapping SHOULD be used. The different prefixes are listed in the TV, the MO is set to "match-mapping" and the CDA is set to "mapping-sent". See Figure 29

Otherwise, the TV contains the prefix, the MO is set to "equal" and the CDA is set to "value-sent".



### 10.7.2. IPv6 source and destination IID

If the DEV or APP IID are based on an LPWAN address, then the IID can be reconstructed with information coming from the LPWAN header. In that case, the TV is not set, the MO is set to "ignore" and the CDA is set to "DevIID" or "AppIID". Note that the LPWAN technology generally carries a single identifier corresponding to the DEV. Therefore AppIID cannot be used.

For privacy reasons or if the DEV address is changing over time, a static value that is not equal to the DEV address SHOULD be used. In that case, the TV contains the static value, the MO operator is set to "equal" and the CDF is set to "not-sent". [RFC7217] provides some methods that MAY be used to derive this static identifier.

If several IIDs are possible, then the TV contains the list of possible IIDs, the MO is set to "match-mapping" and the CDA is set to "mapping-sent".

It MAY also happen that the IID variability only expresses itself on a few bytes. In that case, the TV is set to the stable part of the IID, the MO is set to "MSB" and the CDA is set to "LSB".

Finally, the IID can be sent in extenso on the LPWAN. In that case, the TV is not set, the MO is set to "ignore" and the CDA is set to "value-sent".

### 10.8. IPv6 extensions

No Rule is currently defined that processes IPv6 extensions. If such extensions are needed, their compression/decompression Rules can be based on the MOs and CDAs described above.

### 10.9. UDP source and destination port

To allow for a single Rule being used for both directions, the UDP port values are identified by their role (DEV or APP) and not by their position in the frame (source or destination). The SCHC C/D MUST be aware of the traffic direction (Uplink, Downlink) to select the appropriate field. The following Rules apply for DEV and APP port numbers.

If both ends know the port number, it can be elided. The TV contains the port number, the MO is set to "equal" and the CDA is set to "not-sent".

If the port variation is on few bits, the TV contains the stable part of the port number, the MO is set to "MSB" and the CDA is set to "LSB".

If some well-known values are used, the TV can contain the list of these values, the MO is set to "match-mapping" and the CDA is set to "mapping-sent".

Otherwise the port numbers are sent over the LPWAN. The TV is not set, the MO is set to "ignore" and the CDA is set to "value-sent".

#### 10.10. UDP length field

The UDP length can be computed from the received data. In that case, the TV is not set, the MO is set to "ignore" and the CDA is set to "compute-length".

If the payload is small, the TV can be set to 0x0000, the MO set to "MSB" and the CDA to "LSB".

In other cases, the length SHOULD be sent and the CDA is replaced by "value-sent".

#### 10.11. UDP Checksum field

The UDP checksum operation is mandatory with IPv6 [RFC8200] for most packets but recognizes that there are exceptions to that default behavior.

For instance, protocols that use UDP as a tunnel encapsulation may enable zero-checksum mode for a specific port (or set of ports) for sending and/or receiving. [RFC8200] also stipulates that any node implementing zero-checksum mode must follow the requirements specified in "Applicability Statement for the Use of IPv6 UDP Datagrams with Zero Checksums" [RFC6936].

6LoWPAN Header Compression [RFC6282] also authorizes to send UDP datagram that are deprived of the checksum protection when an upper layer guarantees the integrity of the UDP payload and pseudo-header all the way between the compressor that elides the UDP checksum and the decompressor that computes again it. A specific example of this is when a Message Integrity Check (MIC) protects the compressed message all along that path with a strength that is identical or better to the UDP checksum.

In a similar fashion, this specification allows a SCHC compressor to elide the UDP checks when another layer guarantees an identical or better integrity protection for the UDP payload and the pseudo-

header. In this case, the TV is not set, the MO is set to "ignore" and the CDA is set to "compute-checksum".

In particular, when SCHC fragmentation is used, a fragmentation MIC of 2 bytes or more provides equal or better protection than the UDP checksum; in that case, if the compressor is collocated with the fragmentation point and the decompressor is collocated with the packet reassembly point, then compressor MAY elide the UDP checksum. Whether and when the UDP Checksum is elided is to be specified in the technology-specific documents.

Since the compression happens before the fragmentation, implementors should understand the risks when dealing with unprotected data below the transport layer and take special care when manipulating that data.

In other cases, the checksum SHOULD be explicitly sent. The TV is not set, the MO is set to "ignore" and the CDA is set to "value-sent".

## 11. IANA Considerations

This document has no request to IANA.

## 12. Security considerations

### 12.1. Security considerations for SCHC Compression/Decompression

A malicious header compression could cause the reconstruction of a wrong packet that does not match with the original one. Such a corruption MAY be detected with end-to-end authentication and integrity mechanisms. Header Compression does not add more security problem than what is already needed in a transmission. For instance, to avoid an attack, never re-construct a packet bigger than some configured size (with 1500 bytes as generic default).

### 12.2. Security considerations for SCHC Fragmentation/Reassembly

This subsection describes potential attacks to LPWAN SCHC F/R and suggests possible countermeasures.

A node can perform a buffer reservation attack by sending a first SCHC Fragment to a target. Then, the receiver will reserve buffer space for the IPv6 packet. Other incoming fragmented SCHC Packets will be dropped while the reassembly buffer is occupied during the reassembly timeout. Once that timeout expires, the attacker can repeat the same procedure, and iterate, thus creating a denial of service attack. The (low) cost to mount this attack is linear with

the number of buffers at the target node. However, the cost for an attacker can be increased if individual SCHC Fragments of multiple packets can be stored in the reassembly buffer. To further increase the attack cost, the reassembly buffer can be split into SCHC Fragment-sized buffer slots. Once a packet is complete, it is processed normally. If buffer overload occurs, a receiver can discard packets based on the sender behavior, which MAY help identify which SCHC Fragments have been sent by an attacker.

In another type of attack, the malicious node is required to have overhearing capabilities. If an attacker can overhear a SCHC Fragment, it can send a spoofed duplicate (e.g. with random payload) to the destination. If the LPWAN technology does not support suitable protection (e.g. source authentication and frame counters to prevent replay attacks), a receiver cannot distinguish legitimate from spoofed SCHC Fragments. Therefore, the original IPv6 packet will be considered corrupt and will be dropped. To protect resource-constrained nodes from this attack, it has been proposed to establish a binding among the SCHC Fragments to be transmitted by a node, by applying content-chaining to the different SCHC Fragments, based on cryptographic hash functionality. The aim of this technique is to allow a receiver to identify illegitimate SCHC Fragments.

Further attacks MAY involve sending overlapped fragments (i.e. comprising some overlapping parts of the original IPv6 datagram). Implementers SHOULD make sure that the correct operation is not affected by such event.

In ACK-on-Error, a malicious node MAY force a SCHC Fragment sender to resend a SCHC Fragment a number of times, with the aim to increase consumption of the SCHC Fragment sender's resources. To this end, the malicious node MAY repeatedly send a fake ACK to the SCHC Fragment sender, with a Bitmap that reports that one or more SCHC Fragments have been lost. In order to mitigate this possible attack, MAX\_ACK\_RETRIES MAY be set to a safe value which allows to limit the maximum damage of the attack to an acceptable extent. However, note that a high setting for MAX\_ACK\_RETRIES benefits SCHC Fragment reliability modes, therefore the trade-off needs to be carefully considered.

### 13. Acknowledgements

Thanks to Carsten Bormann, Philippe Clavier, Eduardo Ingles Sanchez, Arunprabhu Kandasamy, Rahul Jadhav, Sergio Lopez Bernal, Antony Markovski, Alexander Pelov, Pascal Thubert, Juan Carlos Zuniga, Diego Dujovne, Edgar Ramos, and Shoichi Sakane for useful design consideration and comments.

## 14. References

### 14.1. 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>>.
- [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>>.
- [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>>.

### 14.2. Informative References

- [RFC3385] Sheinwald, D., Satran, J., Thaler, P., and V. Cavanna, "Internet Protocol Small Computer System Interface (iSCSI) Cyclic Redundancy Check (CRC)/Checksum Considerations", RFC 3385, DOI 10.17487/RFC3385, September 2002, <<https://www.rfc-editor.org/info/rfc3385>>.
- [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>>.
- [RFC5795] Sandlund, K., Pelletier, G., and L-E. Jonsson, "The RObust Header Compression (ROHC) Framework", RFC 5795, DOI 10.17487/RFC5795, March 2010, <<https://www.rfc-editor.org/info/rfc5795>>.
- [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>>.
- [RFC6936] Fairhurst, G. and M. Westerlund, "Applicability Statement for the Use of IPv6 UDP Datagrams with Zero Checksums", RFC 6936, DOI 10.17487/RFC6936, April 2013, <<https://www.rfc-editor.org/info/rfc6936>>.

- [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>>.
- [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>>.
- [RFC8376] Farrell, S., Ed., "Low-Power Wide Area Network (LPWAN) Overview", RFC 8376, DOI 10.17487/RFC8376, May 2018, <<https://www.rfc-editor.org/info/rfc8376>>.

#### Appendix A. SCHC Compression Examples

This section gives some scenarios of the compression mechanism for IPv6/UDP. The goal is to illustrate the behavior of SCHC.

The most common case using the mechanisms defined in this document will be a LPWAN Dev that embeds some applications running over CoAP. In this example, three flows are considered. The first flow is for the device management based on CoAP using Link Local IPv6 addresses and UDP ports 123 and 124 for Dev and App, respectively. The second flow will be a CoAP server for measurements done by the Device (using ports 5683) and Global IPv6 Address prefixes alpha::IID/64 to beta::1/64. The last flow is for legacy applications using different ports numbers, the destination IPv6 address prefix is gamma::1/64.

Figure 28 presents the protocol stack for this Device. IPv6 and UDP are represented with dotted lines since these protocols are compressed on the radio link.




Figure 28: Simplified Protocol Stack for LP-WAN

Note that in some LPWAN technologies, only the Devs have a device ID. Therefore, when such technologies are used, it is necessary to statically define an IID for the Link Local address for the SCHC C/D.

Rule 0

Field	FL	FP	DI	Value	Match Opera.	Comp Decomp Action	Sent [bits]
IPv6 version	4	1	Bi	6	equal	not-sent	
IPv6 DiffServ	8	1	Bi	0	equal	not-sent	
IPv6 Flow Label	20	1	Bi	0	equal	not-sent	
IPv6 Length	16	1	Bi		ignore	comp-length	
IPv6 Next Header	8	1	Bi	17	equal	not-sent	
IPv6 Hop Limit	8	1	Bi	255	ignore	not-sent	
IPv6 DEVprefix	64	1	Bi	FE80::/64	equal	not-sent	
IPv6 DevIID	64	1	Bi		ignore	DevIID	
IPv6 APPprefix	64	1	Bi	FE80::/64	equal	not-sent	
IPv6 AppIID	64	1	Bi	::1	equal	not-sent	
UDP DEVport	16	1	Bi	123	equal	not-sent	
UDP APPport	16	1	Bi	124	equal	not-sent	
UDP Length	16	1	Bi		ignore	comp-length	
UDP checksum	16	1	Bi		ignore	comp-chk	

Rule 1

Field	FL	FP	DI	Value	Match Opera.	Action	Sent [bits]
						Action	

IPv6 version	4	1	Bi	6	equal	not-sent	
IPv6 DiffServ	8	1	Bi	0	equal	not-sent	
IPv6 Flow Label	20	1	Bi	0	equal	not-sent	
IPv6 Length	16	1	Bi		ignore	comp-length	
IPv6 Next Header	8	1	Bi	17	equal	not-sent	
IPv6 Hop Limit	8	1	Bi	255	ignore	not-sent	
IPv6 DEVprefix	64	1	Bi	[alpha/64, fe80::<64]	match- mapping	mapping-sent	[1]
IPv6 DevIID	64	1	Bi		ignore	DevIID	
IPv6 APPprefix	64	1	Bi	[beta/64, alpha/64, fe80::<64]	match- mapping	mapping-sent	[2]
IPv6 AppIID	64	1	Bi	::1000	equal	not-sent	
UDP DEVport	16	1	Bi	5683	equal	not-sent	
UDP APPport	16	1	Bi	5683	equal	not-sent	
UDP Length	16	1	Bi		ignore	comp-length	
UDP checksum	16	1	Bi		ignore	comp-chk	

Rule 2

Field	FL	FP	DI	Value	Match Opera.	Action Action	Sent [bits]
IPv6 version	4	1	Bi	6	equal	not-sent	
IPv6 DiffServ	8	1	Bi	0	equal	not-sent	
IPv6 Flow Label	20	1	Bi	0	equal	not-sent	
IPv6 Length	16	1	Bi		ignore	comp-length	
IPv6 Next Header	8	1	Bi	17	equal	not-sent	
IPv6 Hop Limit	8	1	Up	255	ignore	not-sent	
IPv6 Hop Limit	8	1	Dw		ignore	value-sent	[8]
IPv6 DEVprefix	64	1	Bi	alpha/64	equal	not-sent	
IPv6 DevIID	64	1	Bi		ignore	DevIID	
IPv6 APPprefix	64	1	Bi	gamma/64	equal	not-sent	
IPv6 AppIID	64	1	Bi	::1000	equal	not-sent	
UDP DEVport	16	1	Bi	8720	MSB(12)	LSB	[4]
UDP APPport	16	1	Bi	8720	MSB(12)	LSB	[4]
UDP Length	16	1	Bi		ignore	comp-length	
UDP checksum	16	1	Bi		ignore	comp-chk	

Figure 29: Context Rules



All the fields described in the three Rules depicted on Figure 29 are present in the IPv6 and UDP headers. The DevIID-DID value is found in the L2 header.

The second and third Rules use global addresses. The way the Dev learns the prefix is not in the scope of the document.

The third Rule compresses port numbers to 4 bits.

Appendix B. Fragmentation Examples

This section provides examples for the different fragment reliability modes specified in this document.

Figure 30 illustrates the transmission in No-ACK mode of an IPv6 packet that needs 11 fragments. FCN is 1 bit wide.




Figure 30: Transmission in No-ACK mode of an IPv6 packet carried by 11 fragments

In the following examples, N (i.e. the size of the FCN field) is 3 bits. Therefore, the All-1 FCN value is 7.

Figure 31 illustrates the transmission in ACK-on-Error of an IPv6 packet that needs 11 fragments, with MAX\_WIND\_FCN=6 and no fragment loss.

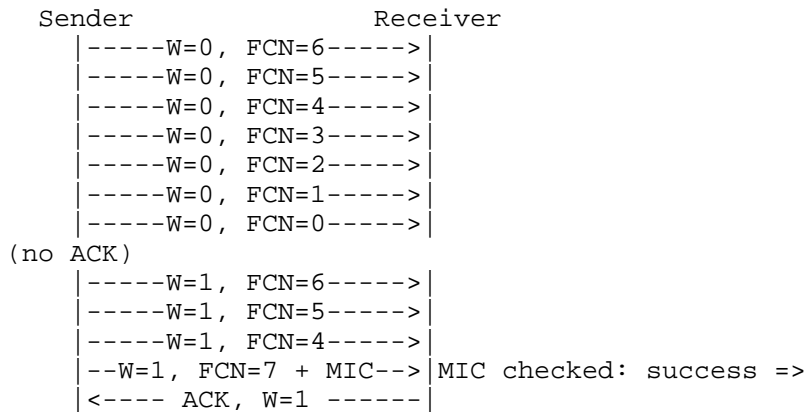


Figure 31: Transmission in ACK-on-Error mode of an IPv6 packet carried by 11 fragments, with MAX\_WIND\_FCN=6 and no loss.

Figure 32 illustrates the transmission in ACK-on-Error mode of an IPv6 packet that needs 11 fragments, with MAX\_WIND\_FCN=6 and three lost fragments.

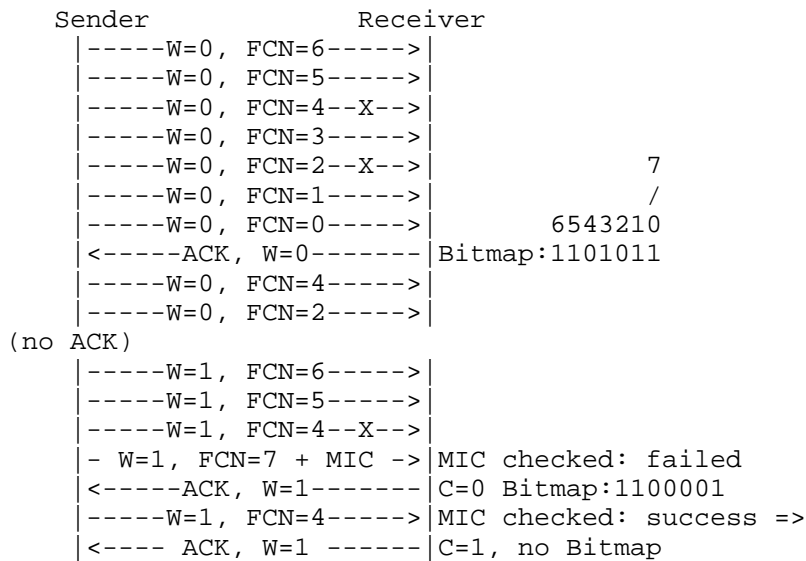


Figure 32: Transmission in ACK-on-Error mode of an IPv6 packet carried by 11 fragments, with MAX\_WIND\_FCN=6 and three lost fragments.

Figure 33 illustrates the transmission in ACK-Always mode of an IPv6 packet that needs 11 fragments, with MAX\_WIND\_FCN=6 and no loss.




Figure 33: Transmission in ACK-Always mode of an IPv6 packet carried by 11 fragments, with MAX\_WIND\_FCN=6 and no lost fragment.

Figure 34 illustrates the transmission in ACK-Always mode of an IPv6 packet that needs 11 fragments, with MAX\_WIND\_FCN=6 and three lost fragments.

```

Sender                                     Receiver
|-----W=1, FCN=6----->|
|-----W=1, FCN=5----->|
|-----W=1, FCN=4--X-->|
|-----W=1, FCN=3----->|
|-----W=1, FCN=2--X-->|                                     7
|-----W=1, FCN=1----->|                                     /
|-----W=1, FCN=0----->|                               6543210
|<-----ACK, W=1----->| Bitmap:1101011
|-----W=1, FCN=4----->|
|-----W=1, FCN=2----->|
|<-----ACK, W=1----->| Bitmap:
|-----W=0, FCN=6----->|
|-----W=0, FCN=5----->|
|-----W=0, FCN=4--X-->|
|--W=0, FCN=7 + MIC-->| MIC checked: failed
|<-----ACK, W=0----->| C= 0 Bitmap:11000001
|-----W=0, FCN=4----->| MIC checked: success =>
|<-----ACK, W=0----->| C= 1 no Bitmap
(End)

```

Figure 34: Transmission in ACK-Always mode of an IPv6 packet carried by 11 fragments, with MAX\_WIND\_FCEN=6 and three lost fragments.

Figure 35 illustrates the transmission in ACK-Always mode of an IPv6 packet that needs 6 fragments, with MAX\_WIND\_FCEN=6, three lost fragments and only one retry needed to recover each lost fragment.

```

Sender                                     Receiver
|-----W=0, FCN=6----->|
|-----W=0, FCN=5----->|
|-----W=0, FCN=4--X-->|
|-----W=0, FCN=3--X-->|
|-----W=0, FCN=2--X-->|
|--W=0, FCN=7 + MIC-->| MIC checked: failed
|<-----ACK, W=0----->| C= 0 Bitmap:11000001
|-----W=0, FCN=4----->| MIC checked: failed
|-----W=0, FCN=3----->| MIC checked: failed
|-----W=0, FCN=2----->| MIC checked: success
|<-----ACK, W=0----->| C=1 no Bitmap
(End)

```

Figure 35: Transmission in ACK-Always mode of an IPv6 packet carried by 11 fragments, with MAX\_WIND\_FCEN=6, three lost fragments and only one retry needed for each lost fragment.

Figure 36 illustrates the transmission in ACK-Always mode of an IPv6 packet that needs 6 fragments, with MAX\_WIND\_FCN=6, three lost fragments, and the second ACK lost.

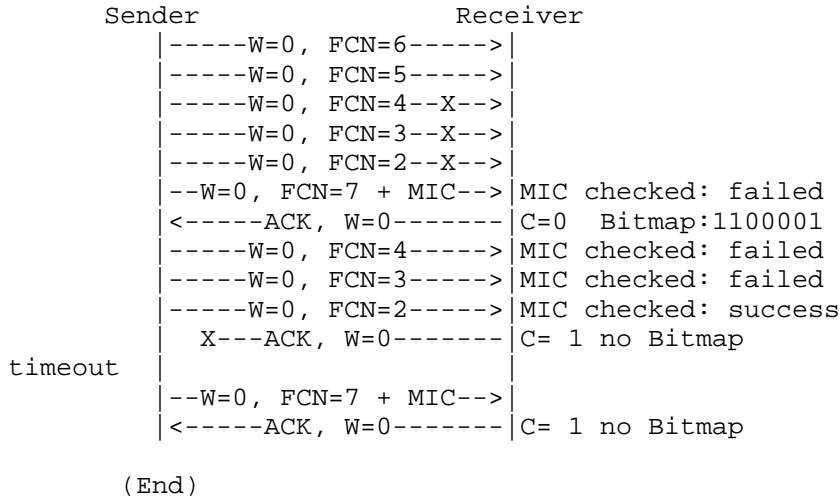


Figure 36: Transmission in ACK-Always mode of an IPv6 packet carried by 11 fragments, with MAX\_WIND\_FCN=6, three lost fragments, and the second ACK lost.

Figure 37 illustrates the transmission in ACK-Always mode of an IPv6 packet that needs 6 fragments, with MAX\_WIND\_FCN=6, with three lost fragments, and one retransmitted fragment lost again.

Sender	Receiver
-----W=0, FCN=6----->	
-----W=0, FCN=5----->	
-----W=0, FCN=4--X-->	
-----W=0, FCN=3--X-->	
-----W=0, FCN=2--X-->	
--W=0, FCN=7 + MIC-->	MIC checked: failed
<-----ACK, W=0-----	C=0 Bitmap:110001
-----W=0, FCN=4----->	MIC checked: failed
-----W=0, FCN=3----->	MIC checked: failed
-----W=0, FCN=2--X-->	
timeout	
--W=0, FCN=7 + MIC-->	All-0 empty
<-----ACK, W=0-----	C=0 Bitmap: 1111101
-----W=0, FCN=2----->	MIC checked: success
<-----ACK, W=0-----	C=1 no Bitmap
(End)	

Figure 37: Transmission in ACK-Always mode of an IPv6 packet carried by 11 fragments, with MAX\_WIND\_FCN=6, with three lost fragments, and one retransmitted fragment lost again.

Figure 38 illustrates the transmission in ACK-Always mode of an IPv6 packet that needs 28 fragments, with N=5, MAX\_WIND\_FCN=23 and two lost fragments. Note that MAX\_WIND\_FCN=23 may be useful when the maximum possible Bitmap size, considering the maximum lower layer technology payload size and the value of R, is 3 bytes. Note also that the FCN of the last fragment of the packet is the one with FCN=31 (i.e.  $FCN=2^N-1$  for N=5, or equivalently, all FCN bits set to 1).

Sender	Receiver
-----W=0, FCN=23----->	
-----W=0, FCN=22----->	
-----W=0, FCN=21--X-->	
-----W=0, FCN=20----->	
-----W=0, FCN=19----->	
-----W=0, FCN=18----->	
-----W=0, FCN=17----->	
-----W=0, FCN=16----->	
-----W=0, FCN=15----->	
-----W=0, FCN=14----->	
-----W=0, FCN=13----->	
-----W=0, FCN=12----->	
-----W=0, FCN=11----->	
-----W=0, FCN=10--X-->	
-----W=0, FCN=9 ----->	
-----W=0, FCN=8 ----->	
-----W=0, FCN=7 ----->	
-----W=0, FCN=6 ----->	
-----W=0, FCN=5 ----->	
-----W=0, FCN=4 ----->	
-----W=0, FCN=3 ----->	
-----W=0, FCN=2 ----->	
-----W=0, FCN=1 ----->	
-----W=0, FCN=0 ----->	
	lcl-Bitmap:11011111111111011111111111
<-----ACK, W=0----->	encoded Bitmap:11011111111111011
-----W=0, FCN=21----->	
-----W=0, FCN=10----->	
<-----ACK, W=0----->	no Bitmap
-----W=1, FCN=23----->	
-----W=1, FCN=22----->	
-----W=1, FCN=21----->	
--W=1, FCN=31 + MIC-->	MIC checked: sucess =>
<-----ACK, W=1----->	no Bitmap

(End)

Figure 38: Transmission in ACK-Always mode of an IPv6 packet carried by 28 fragments, with N=5, MAX\_WIND\_FCN=23 and two lost fragments.

Appendix C. Fragmentation State Machines

The fragmentation state machines of the sender and the receiver, one for each of the different reliability modes, are described in the following figures:

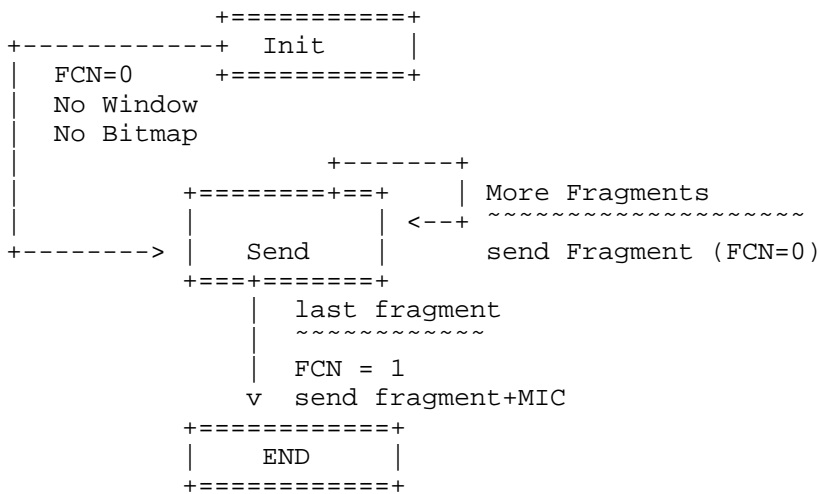


Figure 39: Sender State Machine for the No-ACK Mode

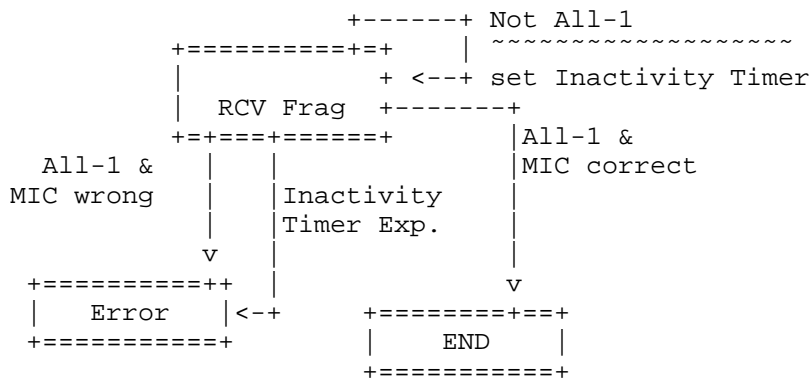


Figure 40: Receiver State Machine for the No-ACK Mode



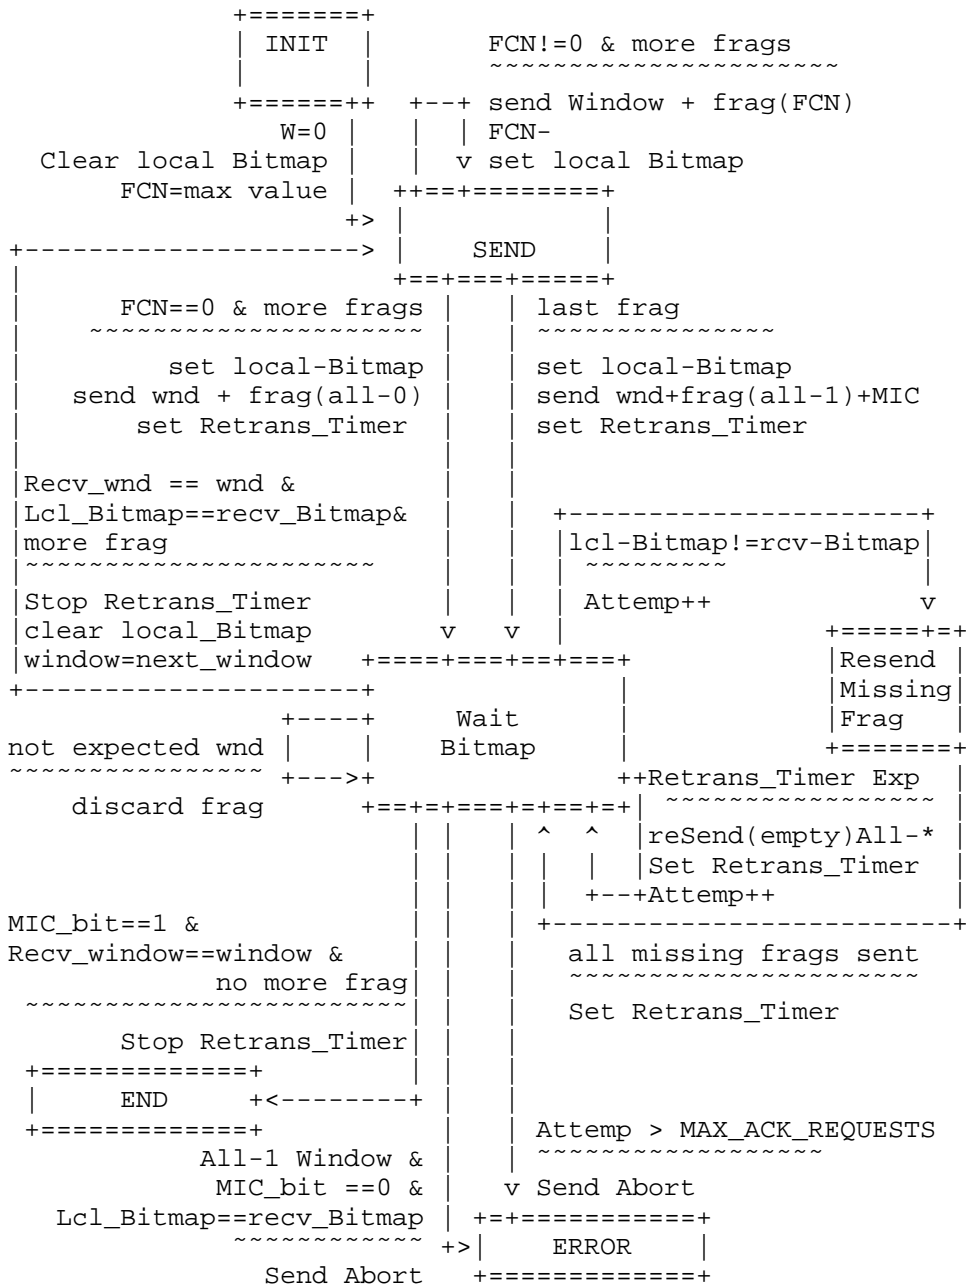



Figure 41: Sender State Machine for the ACK-Always Mode



```
--->* ABORT
~~~~~
      Inactivity_Timer = expires
When DWN_Link
  IF Inactivity_Timer expires
    Send DWL Request
    Attemp++
```

Figure 42: Receiver State Machine for the ACK-Always Mode




Figure 43: Sender State Machine for the ACK-on-Error Mode

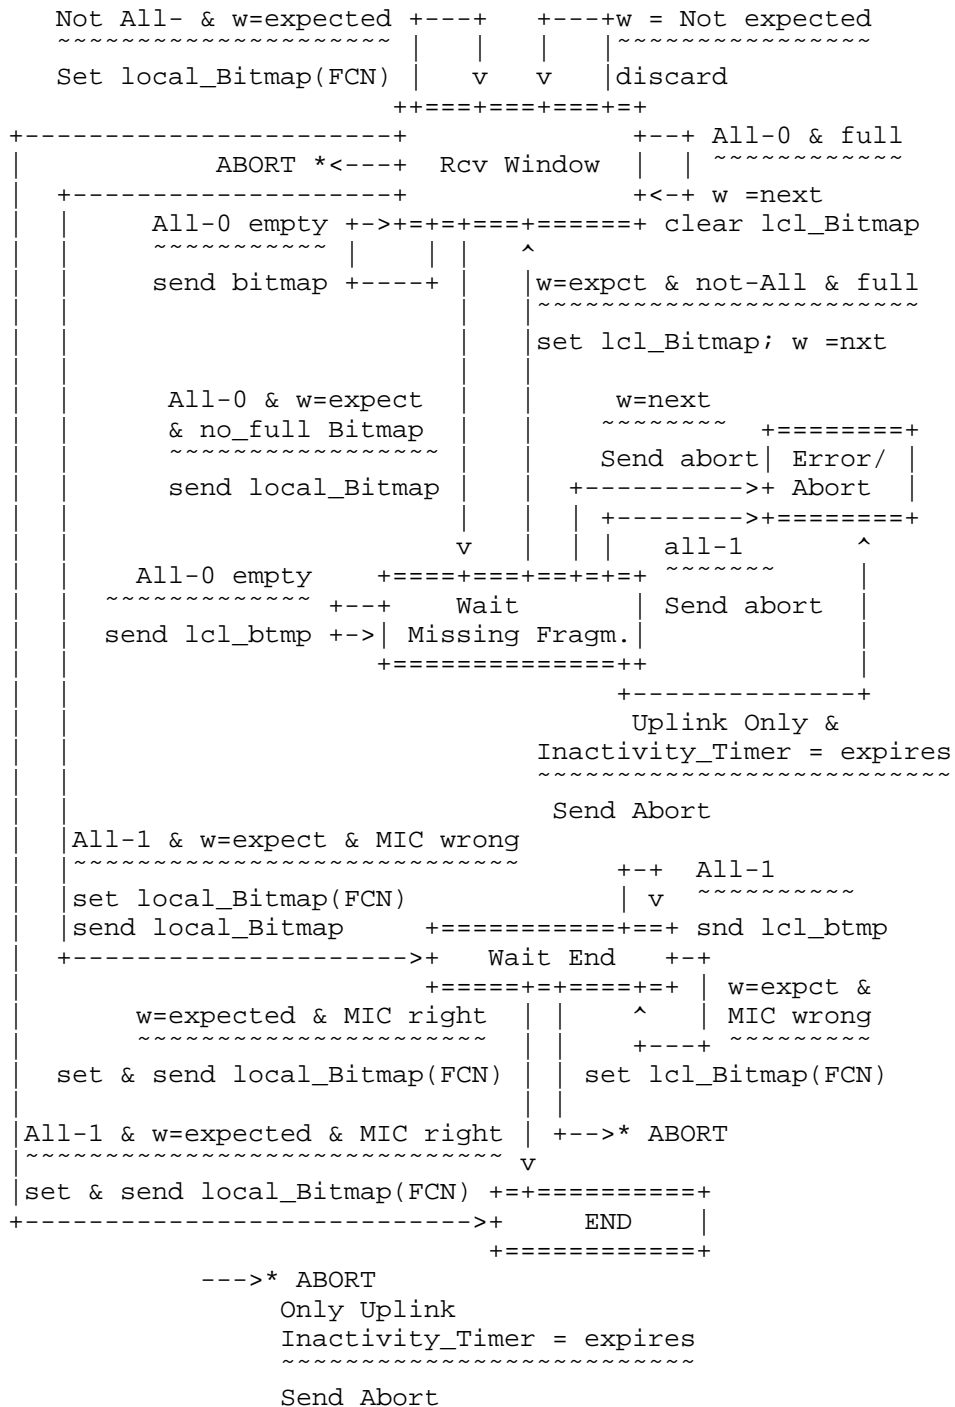


Figure 44: Receiver State Machine for the ACK-on-Error Mode

## Appendix D. SCHC Parameters - Ticket #15

This section gives the list of parameters that need to be defined in the technology-specific documents.

- o Define the most common uses case and how SCHC may be deployed.
- o LPWAN Architecture. Explain the SCHC entities (Compression and Fragmentation), how/where they are represented in the corresponding technology architecture. If applicable, explain the various potential channel conditions for the technology and the corresponding recommended use of C/D and F/R.
- o L2 fragmentation decision
- o Technology developers must evaluate that L2 has strong enough integrity checking to match SCHC's assumption.
- o Rule ID numbering system, number of Rules
- o Size of the Rule IDs
- o The way the Rule ID is sent (L2 or L3) and how (describe)
- o Fragmentation delivery reliability mode used in which cases (e.g. based on link channel condition)
- o Define the number of bits for FCN (N) and DTag (T)
- o in particular, is interleaved packet transmission supported and to what extent
- o The MIC algorithm to be used and the size, if different from the default CRC32
- o Retransmission Timer duration
- o Inactivity Timer duration
- o Define MAX\_ACK\_REQUEST (number of attempts)
- o Padding: size of the L2 Word (for most technologies, a byte; for some technologies, a bit). Value of the padding bits (1 or 0). The value of the padding bits needs to be specified because the padding bits are included in the MIC calculation.

- o Take into account that the length of Rule ID + N + T + W when possible is good to have a multiple of 8 bits to complete a byte and avoid padding
- o In the ACK format to have a length for Rule ID + T + W bit into a complete number of byte to do optimization more easily
- o The technology documents will describe if Rule ID is constrained by any alignment
- o When fragmenting in ACK-on-Error or ACK-Always mode, it is expected that the last window (called All-1 window) will not be fully utilised, i.e. there won't be fragments with all FCN values from MAX\_WIND\_FCN down to 1 and finally All-1. It is worth noting that this document does not mandate that other windows (called All-0 windows) are fully utilised either. This document purposely does not specify that All-1 windows use Bitmaps with the same number of bits as All-0 windows do. By default, Bitmaps for All-0 and All-1 windows are of the same size MAX\_WIND\_FCN + 1. But a technology-specific document MAY revert that decision. The rationale for reverting the decision could be the following: Note that the SCHC ACK sent as a response to an All-1 fragment includes a C bit that SCHC ACK for other windows don't have. Therefore, the SCHC ACK for the All-1 window is one bit bigger. An L2 technology with a severely constrained payload size might decide that this "bump" in the SCHC ACK for the last fragment is a bad resource usage. It could thus mandate that the All-1 window is not allowed to use the FCN value 1 and that the All-1 SCHC ACK Bitmap size is reduced by 1 bit. This provides room for the C bit without creating a bump in the SCHC ACK.

And the following parameters need to be addressed in another document but not forcibly in the technology-specific one:

- o The way the contexts are provisioning
- o The way the Rules are generated

#### Appendix E. Note

Carles Gomez has been funded in part by the Spanish Government (Ministerio de Educacion, Cultura y Deporte) through the Jose Castillejo grant CAS15/00336, and by the ERDF and the Spanish Government through project TEC2016-79988-P. Part of his contribution to this work has been carried out during his stay as a visiting scholar at the Computer Laboratory of the University of Cambridge.

Authors' Addresses

Ana Minaburo  
Acklio  
1137A avenue des Champs Blancs  
35510 Cesson-Sevigne Cedex  
France

Email: ana@ackl.io

Laurent Toutain  
IMT-Atlantique  
2 rue de la Chataigneraie  
CS 17607  
35576 Cesson-Sevigne Cedex  
France

Email: Laurent.Toutain@imt-atlantique.fr

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

Email: carlesgo@entel.upc.edu

Dominique Barthel  
Orange Labs  
28 chemin du Vieux Chene  
38243 Meylan  
France

Email: dominique.barthel@orange.com



lpwan  
Internet-Draft  
Intended status: Informational  
Expires: August 11, 2018

S. Farrell, Ed.  
Trinity College Dublin  
February 7, 2018

LPWAN Overview  
draft-ietf-lpwan-overview-10

Abstract

Low Power Wide Area Networks (LPWAN) are wireless technologies with characteristics such as large coverage areas, low bandwidth, possibly very small packet and application layer data sizes and long battery life operation. This memo is an informational overview of the set of LPWAN technologies being considered in the IETF and of the gaps that exist between the needs of those technologies and the goal of running IP in LPWANs.

Status of This Memo

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

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

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

This Internet-Draft will expire on August 11, 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 (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of

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

Table of Contents

- 1. Introduction . . . . . 3
- 2. LPWAN Technologies . . . . . 3
  - 2.1. LoRaWAN . . . . . 4
    - 2.1.1. Provenance and Documents . . . . . 4
    - 2.1.2. Characteristics . . . . . 4
  - 2.2. Narrowband IoT (NB-IoT) . . . . . 11
    - 2.2.1. Provenance and Documents . . . . . 11
    - 2.2.2. Characteristics . . . . . 11
  - 2.3. SIGFOX . . . . . 15
    - 2.3.1. Provenance and Documents . . . . . 15
    - 2.3.2. Characteristics . . . . . 16
  - 2.4. Wi-SUN Alliance Field Area Network (FAN) . . . . . 20
    - 2.4.1. Provenance and Documents . . . . . 20
    - 2.4.2. Characteristics . . . . . 21
- 3. Generic Terminology . . . . . 24
- 4. Gap Analysis . . . . . 25
  - 4.1. Naive application of IPv6 . . . . . 26
  - 4.2. 6LOWPAN . . . . . 26
    - 4.2.1. Header Compression . . . . . 27
    - 4.2.2. Address Autoconfiguration . . . . . 27
    - 4.2.3. Fragmentation . . . . . 27
    - 4.2.4. Neighbor Discovery . . . . . 28
  - 4.3. 6lo . . . . . 29
  - 4.4. 6tisch . . . . . 29
  - 4.5. RoHC . . . . . 29
  - 4.6. ROLL . . . . . 30
  - 4.7. CoAP . . . . . 30
  - 4.8. Mobility . . . . . 30
  - 4.9. DNS and LPWAN . . . . . 31
- 5. Security Considerations . . . . . 31
- 6. IANA Considerations . . . . . 32
- 7. Contributors . . . . . 32
- 8. Acknowledgments . . . . . 35
- 9. Informative References . . . . . 35
- Appendix A. Changes . . . . . 41
  - A.1. From -00 to -01 . . . . . 41
  - A.2. From -01 to -02 . . . . . 41
  - A.3. From -02 to -03 . . . . . 41
  - A.4. From -03 to -04 . . . . . 42
  - A.5. From -04 to -05 . . . . . 42
  - A.6. From -05 to -06 . . . . . 42
  - A.7. From -06 to -07 . . . . . 42
  - A.8. From -07 to -08 . . . . . 42

A.9. From -08 to -09 . . . . . 43  
A.10. From -09 to -10 . . . . . 43  
Author's Address . . . . . 43

1. Introduction

This document provides background material and an overview of the technologies being considered in the IETF's Low Power Wide-Area Networking (LPWAN) working group. We also provide a gap analysis between the needs of these technologies and currently available IETF specifications.

Most technologies in this space aim for similar goals of supporting large numbers of very low-cost, low-throughput devices with very-low power consumption, so that even battery-powered devices can be deployed for years. LPWAN devices also tend to be constrained in their use of bandwidth, for example with limited frequencies being allowed to be used within limited duty-cycles (usually expressed as a percentage of time per-hour that the device is allowed to transmit.) And as the name implies, coverage of large areas is also a common goal. So, by and large, the different technologies aim for deployment in very similar circumstances.

What mainly distinguishes LPWANs from other constrained networks is that in LPWANs the balancing act related to power consumption/battery life, cost and bandwidth tends to prioritise doing better with respect to power and cost and we are more willing to live with extremely low bandwidth and constrained duty-cycles when making the various trade-offs required, in order to get the multiple-kilometre radio links implied by the "wide area" aspect of the LPWAN term.

Existing pilot deployments have shown huge potential and created much industrial interest in these technologies. As of today, essentially no LPWAN end-devices (other than for Wi-SUN) have IP capabilities. Connecting LPWANs to the Internet would provide significant benefits to these networks in terms of interoperability, application deployment, and management, among others. The goal of the IETF LPWAN working group is to, where necessary, adapt IETF-defined protocols, addressing schemes and naming to this particular constrained environment.

This document is largely the work of the people listed in Section 7.

2. LPWAN Technologies

This section provides an overview of the set of LPWAN technologies that are being considered in the LPWAN working group. The text for each was mainly contributed by proponents of each technology.

Note that this text is not intended to be normative in any sense, but simply to help the reader in finding the relevant layer 2 specifications and in understanding how those integrate with IETF-defined technologies. Similarly, there is no attempt here to set out the pros and cons of the relevant technologies.

Note that some of the technology-specific drafts referenced below may have been updated since publication of this document.

## 2.1. LoRaWAN

### 2.1.1. Provenance and Documents

LoRaWAN is an ISM-based wireless technology for long-range low-power low-data-rate applications developed by the LoRa Alliance, a membership consortium. <https://www.lora-alliance.org/> This draft is based on version 1.0.2 [LoRaSpec] of the LoRa specification. That specification is publicly available and has already seen several deployments across the globe.

### 2.1.2. Characteristics

LoRaWAN aims to support end-devices operating on a single battery for an extended period of time (e.g., 10 years or more), extended coverage through 155 dB maximum coupling loss, and reliable and efficient file download (as needed for remote software/firmware upgrade).

LoRaWAN networks are typically organized in a star-of-stars topology in which gateways relay messages between end-devices and a central "network server" in the backend. Gateways are connected to the network server via IP links while end-devices use single-hop LoRaWAN communication that can be received at one or more gateways. Communication is generally bi-directional; uplink communication from end-devices to the network server is favored in terms of overall bandwidth availability.

Figure 1 shows the entities involved in a LoRaWAN network.




Figure 1: LoRaWAN architecture

- o End-device: a LoRa client device, sometimes called a mote. Communicates with gateways.
- o Gateway: a radio on the infrastructure-side, sometimes called a concentrator or base-station. Communicates with end-devices and, via IP, with a network server.
- o Network Server: The Network Server (NS) terminates the LoRaWAN MAC layer for the end-devices connected to the network. It is the center of the star topology.
- o Join Server: The Join Server (JS) is a server on the Internet side of an NS that processes join requests from an end-devices.
- o Uplink message: refers to communications from an end-device to a network server or application via one or more gateways.
- o Downlink message: refers to communications from a network server or application via one gateway to a single end-device or a group of end-devices (considering multicasting).
- o Application: refers to application layer code both on the end-device and running "behind" the network server. For LoRaWAN, there will generally only be one application running on most end-devices. Interfaces between the network server and application are not further described here.

In LoRaWAN networks, end-device transmissions may be received at multiple gateways, so during nominal operation a network server may see multiple instances of the same uplink message from an end-device.

The LoRaWAN network infrastructure manages the data rate and RF output power for each end-device individually by means of an adaptive data rate (ADR) scheme. End-devices may transmit on any channel allowed by local regulation at any time.

LoRaWAN radios make use of industrial, scientific and medical (ISM) bands, for example, 433MHz and 868MHz within the European Union and 915MHz in the Americas.

The end-device changes channel in a pseudo-random fashion for every transmission to help make the system more robust to interference and/or to conform to local regulations.

Figure 2 below shows that after a transmission slot a Class A device turns on its receiver for two short receive windows that are offset from the end of the transmission window. End-devices can only transmit a subsequent uplink frame after the end of the associated receive windows. When a device joins a LoRaWAN network, there are similar timeouts on parts of that process.




Figure 2: LoRaWAN Class A transmission and reception window

Given the different regional requirements the detailed specification for the LoRaWAN physical layer (taking up more than 30 pages of the specification) is not reproduced here. Instead and mainly to illustrate the kinds of issue encountered, in Table 1 we present some of the default settings for one ISM band (without fully explaining those here) and in Table 2 we describe maxima and minima for some parameters of interest to those defining ways to use IETF protocols over the LoRaWAN MAC layer.

Parameters	Default Value
Rx delay 1	1 s
Rx delay 2	2 s (must be RECEIVE_DELAY1 + 1s)
join delay 1	5 s
join delay 2	6 s
868MHz Default channels	3 (868.1,868.2,868.3), data rate: 0.3-50kbps

Table 1: Default settings for EU 868MHz band

Parameter/Notes	Min	Max
Duty Cycle: some but not all ISM bands impose a limit in terms of how often an end-device can transmit. In some cases LoRaWAN is more restrictive in an attempt to avoid congestion.	1%	no-limit
EU 868MHz band data rate/frame-size	250 bits/s : 59 octets	50000 bits/s : 250 octets
US 915MHz band data rate/frame-size	980 bits/s : 19 octets	21900 bits/s : 250 octets

Table 2: Minima and Maxima for various LoRaWAN Parameters

Note that in the case of the smallest frame size (19 octets), 8 octets are required for LoRa MAC layer headers leaving only 11 octets for payload (including MAC layer options). However, those settings do not apply for the join procedure - end-devices are required to use a channel and data rate that can send the 23-byte Join-request message for the join procedure.

Uplink and downlink higher layer data is carried in a MACPayload. There is a concept of "ports" (an optional 8-bit value) to handle

different applications on an end-device. Port zero is reserved for LoRaWAN specific messaging, such as the configuration of the end device's network parameters (available channels, data rates, ADR parameters, RX1/2 delay, etc.).

In addition to carrying higher layer PDUs there are Join-Request and Join-Response (aka Join-Accept) messages for handling network access. And so-called "MAC commands" (see below) up to 15 bytes long can be piggybacked in an options field ("FOpts").

There are a number of MAC commands for link and device status checking, ADR and duty-cycle negotiation, managing the RX windows and radio channel settings. For example, the link check response message allows the network server (in response to a request from an end-device) to inform an end-device about the signal attenuation seen most recently at a gateway, and to also tell the end-device how many gateways received the corresponding link request MAC command.

Some MAC commands are initiated by the network server. For example, one command allows the network server to ask an end-device to reduce its duty-cycle to only use a proportion of the maximum allowed in a region. Another allows the network server to query the end-device's power status with the response from the end-device specifying whether it has an external power source or is battery powered (in which case a relative battery level is also sent to the network server).

In order to operate nominally on a LoRaWAN network, a device needs a 32-bit device address, that is assigned when the device "joins" the network (see below for the join procedure) or that is pre-provisioned into the device. In case of roaming devices, the device address is assigned based on the 24-bit network identifier (NetID) that is allocated to the network by the LoRa Alliance. Non-roaming devices can be assigned device addresses by the network without relying on a LoRa Alliance-assigned NetID.

End-devices are assumed to work with one or a quite limited number of applications, identified by a 64-bit AppEUI, which is assumed to be a registered IEEE EUI64 value. In addition, a device needs to have two symmetric session keys, one for protecting network artifacts (port=0), the NwkSKey, and another for protecting application layer traffic, the AppSKey. Both keys are used for 128-bit AES cryptographic operations. So, one option is for an end-device to have all of the above, plus channel information, somehow (pre-)provisioned, in which case the end-device can simply start transmitting. This is achievable in many cases via out-of-band means given the nature of LoRaWAN networks. Table 3 summarizes these values.



Value	Description
DevAddr	DevAddr (32-bits) = device-specific network address generated from the NetID
AppEUI	IEEE EUI64 corresponding to the join server for an application
NwksKey	128-bit network session key used with AES-CMAC
AppSKey	128-bit application session key used with AES-CTR
AppKey	128-bit application session key used with AES-ECB

Table 3: Values required for nominal operation

As an alternative, end-devices can use the LoRaWAN join procedure with a join server behind the NS in order to setup some of these values and dynamically gain access to the network. To use the join procedure, an end-device must still know the AppEUI, and in addition, a different (long-term) symmetric key that is bound to the AppEUI - this is the application key (AppKey), and is distinct from the application session key (AppSKey). The AppKey is required to be specific to the device, that is, each end-device should have a different AppKey value. And finally, the end-device also needs a long-term identifier for itself, syntactically also an EUI-64, and known as the device EUI or DevEUI. Table 4 summarizes these values.

Value	Description
DevEUI	IEEE EUI64 naming the device
AppEUI	IEEE EUI64 naming the application
AppKey	128-bit long term application key for use with AES

Table 4: Values required for join procedure

The join procedure involves a special exchange where the end-device asserts the AppEUI and DevEUI (integrity protected with the long-term AppKey, but not encrypted) in a Join-request uplink message. This is then routed to the network server which interacts with an entity that knows that AppKey to verify the Join-request. All going well, a Join-accept downlink message is returned from the network server to

the end-device that specifies the 24-bit NetID, 32-bit DevAddr and channel information and from which the AppSKey and NwkSKey can be derived based on knowledge of the AppKey. This provides the end-device with all the values listed in Table 3.

All payloads are encrypted and have data integrity. MAC commands, when sent as a payload (port zero), are therefore protected. MAC commands piggy-backed as frame options ("FOpts") are however sent in clear. Any MAC commands sent as frame options and not only as payload, are visible to a passive attacker but are not malleable for an active attacker due to the use of the Message Integrity Check (MIC) described below.

For LoRaWAN version 1.0.x, the NwkSKey session key is used to provide data integrity between the end-device and the network server. The AppSKey is used to provide data confidentiality between the end-device and network server, or to the application "behind" the network server, depending on the implementation of the network.

All MAC layer messages have an outer 32-bit MIC calculated using AES-CMAC calculated over the ciphertext payload and other headers and using the NwkSKey. Payloads are encrypted using AES-128, with a counter-mode derived from IEEE 802.15.4 using the AppSKey. Gateways are not expected to be provided with the AppSKey or NwkSKey, all of the infrastructure-side cryptography happens in (or "behind") the network server. When session keys are derived from the AppKey as a result of the join procedure the Join-accept message payload is specially handled.

The long-term AppKey is directly used to protect the Join-accept message content, but the function used is not an AES-encrypt operation, but rather an AES-decrypt operation. The justification is that this means that the end-device only needs to implement the AES-encrypt operation. (The counter mode variant used for payload decryption means the end-device doesn't need an AES-decrypt primitive.)

The Join-accept plaintext is always less than 16 bytes long, so electronic code book (ECB) mode is used for protecting Join-accept messages. The Join-accept contains an AppNonce (a 24 bit value) that is recovered on the end-device along with the other Join-accept content (e.g. DevAddr) using the AES-encrypt operation. Once the Join-accept payload is available to the end-device the session keys are derived from the AppKey, AppNonce and other values, again using an ECB mode AES-encrypt operation, with the plaintext input being a maximum of 16 octets.

## 2.2. Narrowband IoT (NB-IoT)

### 2.2.1. Provenance and Documents

Narrowband Internet of Things (NB-IoT) is developed and standardized by 3GPP. The standardization of NB-IoT was finalized with 3GPP Release 13 in June 2016, and further enhancements for NB-IoT are specified in 3GPP Release 14 in 2017, for example in the form of multicast support. Further features and improvements will be developed in the following releases, but NB-IoT has been ready to be deployed since 2016, and is rather simple to deploy especially in the existing LTE networks with a software upgrade in the operator's base stations. For more information of what has been specified for NB-IoT, 3GPP specification 36.300 [TGPP36300] provides an overview and overall description of the E-UTRAN radio interface protocol architecture, while specifications 36.321 [TGPP36321], 36.322 [TGPP36322], 36.323 [TGPP36323] and 36.331 [TGPP36331] give more detailed description of MAC, Radio Link Control (RLC), Packet Data Convergence Protocol (PDCP) and Radio Resource Control (RRC) protocol layers, respectively. Note that the description below assumes familiarity with numerous 3GPP terms.

For a general overview of NB-IoT, see [nbiot-ov].

### 2.2.2. Characteristics

Specific targets for NB-IoT include: Less than US\$5 module cost, extended coverage of 164 dB maximum coupling loss, battery life of over 10 years, ~55000 devices per cell and uplink reporting latency of less than 10 seconds.

NB-IoT supports Half Duplex FDD operation mode with 60 kbps peak rate in uplink and 30 kbps peak rate in downlink, and a maximum transmission unit (MTU) size of 1600 bytes limited by PDCP layer (see Figure 4 for the protocol structure), which is the highest layer in the user plane, as explained later. Any packet size up to the said MTU size can be passed to the NB-IoT stack from higher layers, segmentation of the packet is performed in the RLC layer, which can segment the data to transmission blocks with size as small as 16 bits. As the name suggests, NB-IoT uses narrowbands with bandwidth of 180 kHz in both downlink and uplink. The multiple access scheme used in the downlink is OFDMA with 15 kHz sub-carrier spacing. In uplink, SC-FDMA single tone with either 15kHz or 3.75 kHz tone spacing is used, or optionally multi-tone SC-FDMA can be used with 15 kHz tone spacing.

NB-IoT can be deployed in three ways. In-band deployment means that the narrowband is deployed inside the LTE band and radio resources

are flexibly shared between NB-IoT and normal LTE carrier. In Guard-band deployment the narrowband uses the unused resource blocks between two adjacent LTE carriers. Standalone deployment is also supported, where the narrowband can be located alone in dedicated spectrum, which makes it possible for example to reframe a GSM carrier at 850/900 MHz for NB-IoT. All three deployment modes are used in licensed frequency bands. The maximum transmission power is either 20 or 23 dBm for uplink transmissions, while for downlink transmission the eNodeB may use higher transmission power, up to 46 dBm depending on the deployment.

A maximum coupling loss (MCL) target for NB-IoT coverage enhancements defined by 3GPP is 164 dB. With this MCL, the performance of NB-IoT in downlink varies between 200 bps and 2-3 kbps, depending on the deployment mode. Stand-alone operation may achieve the highest data rates, up to few kbps, while in-band and guard-band operations may reach several hundreds of bps. NB-IoT may even operate with MCL higher than 170 dB with very low bit rates.

For signaling optimization, two options are introduced in addition to legacy LTE RRC connection setup; mandatory Data-over-NAS (Control Plane optimization, solution 2 in [TGPP23720]) and optional RRC Suspend/Resume (User Plane optimization, solution 18 in [TGPP23720]). In the control plane optimization the data is sent over Non-Access Stratum, directly to/from Mobility Management Entity (MME) (see Figure 3 for the network architecture) in the core network to the User Equipment (UE) without interaction from the base station. This means there are no Access Stratum security or header compression provided by the PDCP layer in the eNodeB, as the Access Stratum is bypassed, and only limited RRC procedures. RoHC based header compression may still optionally be provided and terminated in MME.

The RRC Suspend/Resume procedures reduce the signaling overhead required for UE state transition from RRC Idle to RRC Connected mode compared to legacy LTE operation in order to have quicker user plane transaction with the network and return to RRC Idle mode faster.

In order to prolong device battery life, both power-saving mode (PSM) and extended DRX (eDRX) are available to NB-IoT. With eDRX the RRC Connected mode DRX cycle is up to 10.24 seconds and in RRC Idle the eDRX cycle can be up to 3 hours. In PSM the device is in a deep sleep state and only wakes up for uplink reporting, after which there is a window, configured by the network, during which the device receiver is open for downlink connectivity, or for periodical "keep-alive" signaling (PSM uses periodic TAU signaling with additional reception window for downlink reachability).

Since NB-IoT operates in licensed spectrum, it has no channel access restrictions allowing up to a 100% duty-cycle.

3GPP access security is specified in [TGPP33203].

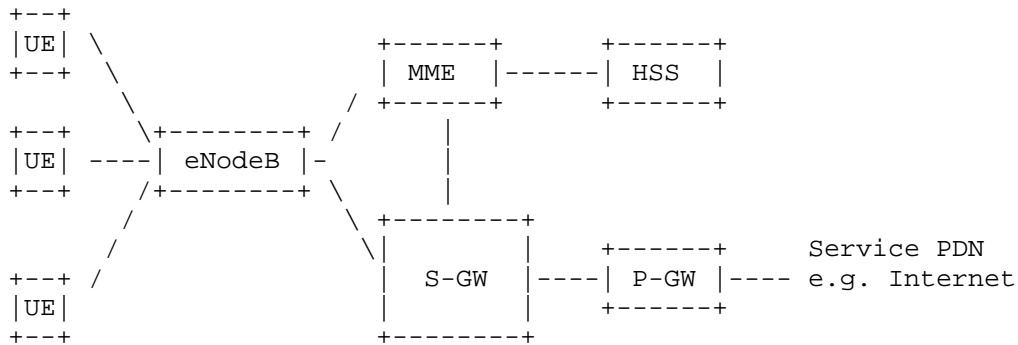


Figure 3: 3GPP network architecture

Figure 3 shows the 3GPP network architecture, which applies to NB-IoT. Mobility Management Entity (MME) is responsible for handling the mobility of the UE. MME tasks include tracking and paging UEs, session management, choosing the Serving gateway for the UE during initial attachment and authenticating the user. At MME, the Non-Access Stratum (NAS) signaling from the UE is terminated.

Serving Gateway (S-GW) routes and forwards the user data packets through the access network and acts as a mobility anchor for UEs during handover between base stations known as eNodeBs and also during handovers between NB-IoT and other 3GPP technologies.

Packet Data Network Gateway (P-GW) works as an interface between 3GPP network and external networks.

The Home Subscriber Server (HSS) contains user-related and subscription-related information. It is a database, which performs mobility management, session establishment support, user authentication and access authorization.

E-UTRAN consists of components of a single type, eNodeB. eNodeB is a base station, which controls the UEs in one or several cells.

The 3GPP radio protocol architecture is illustrated in Figure 4.




Figure 4: 3GPP radio protocol architecture for control plane

Control plane protocol stack

The radio protocol architecture of NB-IoT (and LTE) is separated into control plane and user plane. The control plane consists of protocols which control the radio access bearers and the connection between the UE and the network. The highest layer of control plane is called Non-Access Stratum (NAS), which conveys the radio signaling between the UE and the Evolved Packet Core (EPC), passing transparently through the radio network. NAS responsible for authentication, security control, mobility management and bearer management.

Access Stratum (AS) is the functional layer below NAS, and in the control plane it consists of Radio Resource Control protocol (RRC) [TGPP36331], which handles connection establishment and release functions, broadcast of system information, radio bearer establishment, reconfiguration and release. RRC configures the user and control planes according to the network status. There exists two RRC states, RRC\_Idle or RRC\_Connected, and RRC entity controls the switching between these states. In RRC\_Idle, the network knows that the UE is present in the network and the UE can be reached in case of incoming call/downlink data. In this state, the UE monitors paging, performs cell measurements and cell selection and acquires system information. Also the UE can receive broadcast and multicast data, but it is not expected to transmit or receive unicast data. In RRC\_Connected the UE has a connection to the eNodeB, the network knows the UE location on the cell level and the UE may receive and transmit unicast data. An RRC connection is established when the UE is expected to be active in the network, to transmit or receive data. The RRC connection is released, switching back to RRC\_Idle, when

there is no more traffic in order to preserve UE battery life and radio resources. However, a new feature was introduced for NB-IoT, as mentioned earlier, which allows data to be transmitted from the MME directly to the UE transparently to the eNodeB, thus bypassing AS functions.

Packet Data Convergence Protocol's (PDCP) [TGPP36323] main services in control plane are transfer of control plane data, ciphering and integrity protection.

Radio Link Control protocol (RLC) [TGPP36322] performs transfer of upper layer PDUs and optionally error correction with Automatic Repeat reQuest (ARQ), concatenation, segmentation, and reassembly of RLC SDUs, in-sequence delivery of upper layer PDUs, duplicate detection, RLC SDU discard, RLC-re-establishment and protocol error detection and recovery.

Medium Access Control protocol (MAC) [TGPP36321] provides mapping between logical channels and transport channels, multiplexing of MAC SDUs, scheduling information reporting, error correction with HARQ, priority handling and transport format selection.

Physical layer [TGPP36201] provides data transport services to higher layers. These include error detection and indication to higher layers, FEC encoding, HARQ soft-combining, rate matching and mapping of the transport channels onto physical channels, power weighting and modulation of physical channels, frequency and time synchronization and radio characteristics measurements.

User plane is responsible for transferring the user data through the Access Stratum. It interfaces with IP and the highest layer of user plane is PDCP, which in user plane performs header compression using Robust Header Compression (RoHC), transfer of user plane data between eNodeB and UE, ciphering and integrity protection. Similar to control plane, lower layers in user plane include RLC, MAC and physical layer performing the same tasks as in control plane.

## 2.3. SIGFOX

### 2.3.1. Provenance and Documents

The SIGFOX LPWAN is in line with the terminology and specifications being defined by ETSI [etsi\_unb]. As of today, SIGFOX's network has been fully deployed in 12 countries, with ongoing deployments on 26 other countries, giving in total a geography of 2 million square kilometers, containing 512 million people.

### 2.3.2. Characteristics

SIGFOX LPWAN autonomous battery-operated devices send only a few bytes per day, week or month, in principle allowing them to remain on a single battery for up to 10-15 years. Hence, the system is designed as to allow devices to last several years, sometimes even buried underground.

Since the radio protocol is connection-less and optimized for uplink communications, the capacity of a SIGFOX base station depends on the number of messages generated by devices, and not on the actual number of devices. Likewise, the battery life of devices depends on the number of messages generated by the device. Depending on the use case, devices can vary from sending less than one message per device per day, to dozens of messages per device per day.

The coverage of the cell depends on the link budget and on the type of deployment (urban, rural, etc.). The radio interface is compliant with the following regulations:

Spectrum allocation in the USA [fcc\_ref]

Spectrum allocation in Europe [etsi\_ref]

Spectrum allocation in Japan [arib\_ref]

The SIGFOX radio interface is also compliant with the local regulations of the following countries: Australia, Brazil, Canada, Kenya, Lebanon, Mauritius, Mexico, New Zealand, Oman, Peru, Singapore, South Africa, South Korea, and Thailand.

The radio interface is based on Ultra Narrow Band (UNB) communications, which allow an increased transmission range by spending a limited amount of energy at the device. Moreover, UNB allows a large number of devices to coexist in a given cell without significantly increasing the spectrum interference.

Both uplink and downlink are supported, although the system is optimized for uplink communications. Due to spectrum optimizations, different uplink and downlink frames and time synchronization methods are needed.

The main radio characteristics of the UNB uplink transmission are:

- o Channelization mask: 100 Hz / 600 Hz (depending on the region)
- o Uplink baud rate: 100 baud / 600 baud (depending on the region)



- o Modulation scheme: DBPSK
- o Uplink transmission power: compliant with local regulation
- o Link budget: 155 dB (or better)
- o Central frequency accuracy: not relevant, provided there is no significant frequency drift within an uplink packet transmission

For example, in Europe the UNB uplink frequency band is limited to 868.00 to 868.60 MHz, with a maximum output power of 25 mW and a duty cycle of 1%.

The format of the uplink frame is the following:




Figure 5: Uplink Frame Format

The uplink frame is composed of the following fields:

- o Preamble: 19 bits
- o Frame sync and header: 29 bits
- o Device ID: 32 bits
- o Payload: 0-96 bits
- o Authentication: 16-40 bits
- o Frame check sequence: 16 bits (CRC)

The main radio characteristics of the UNB downlink transmission are:

- o Channelization mask: 1.5 kHz
- o Downlink baud rate: 600 baud
- o Modulation scheme: GFSK
- o Downlink transmission power: 500 mW / 4W (depending on the region)
- o Link budget: 153 dB (or better)

- o Central frequency accuracy: the center frequency of downlink transmission is set by the network according to the corresponding uplink transmission

For example, in Europe the UNB downlink frequency band is limited to 869.40 to 869.65 MHz, with a maximum output power of 500 mW with 10% duty cycle.

The format of the downlink frame is the following:




Figure 6: Downlink Frame Format

The downlink frame is composed of the following fields:

- o Preamble: 91 bits
- o Frame sync and header: 13 bits
- o Error Correcting Code (ECC): 32 bits
- o Payload: 0-64 bits
- o Authentication: 16 bits
- o Frame check sequence: 8 bits (CRC)

The radio interface is optimized for uplink transmissions, which are asynchronous. Downlink communications are achieved by devices querying the network for available data.

A device willing to receive downlink messages opens a fixed window for reception after sending an uplink transmission. The delay and duration of this window have fixed values. The network transmits the downlink message for a given device during the reception window, and the network also selects the base station (BS) for transmitting the corresponding downlink message.

Uplink and downlink transmissions are unbalanced due to the regulatory constraints on ISM bands. Under the strictest regulations, the system can allow a maximum of 140 uplink messages and 4 downlink messages per device per day. These restrictions can

be slightly relaxed depending on system conditions and the specific regulatory domain of operation.

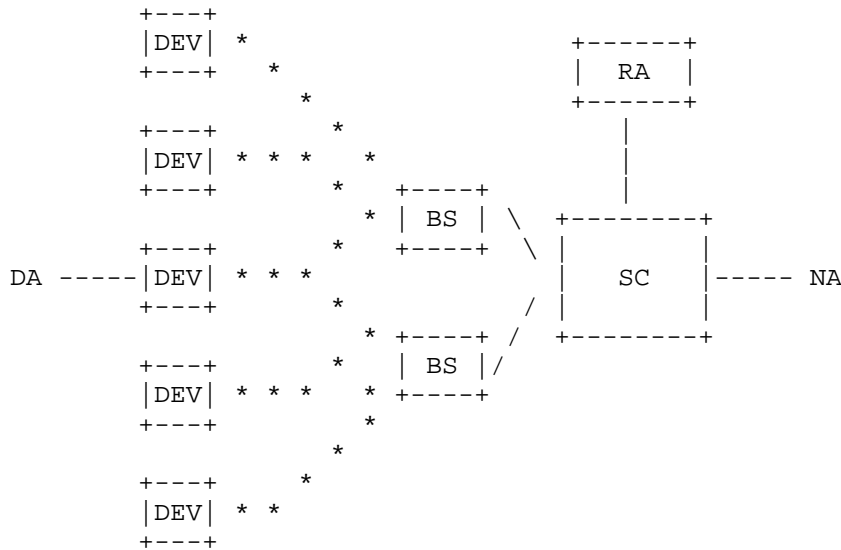


Figure 7: SIGFOX network architecture

Figure 7 depicts the different elements of the SIGFOX network architecture.

SIGFOX has a "one-contract one-network" model allowing devices to connect in any country, without any need or notion of either roaming or handover.

The architecture consists of a single cloud-based core network, which allows global connectivity with minimal impact on the end device and radio access network. The core network elements are the Service Center (SC) and the Registration Authority (RA). The SC is in charge of the data connectivity between the Base Station (BS) and the Internet, as well as the control and management of the BSs and End Points. The RA is in charge of the End Point network access authorization.

The radio access network is comprised of several BSs connected directly to the SC. Each BS performs complex L1/L2 functions, leaving some L2 and L3 functionalities to the SC.

The Devices (DEVs) or End Points (EPs) are the objects that communicate application data between local device applications (DAs) and network applications (NAs).

Devices (or EPs) can be static or nomadic, as they associate with the SC and they do not attach to any specific BS. Hence, they can communicate with the SC through one or multiple BSs.

Due to constraints in the complexity of the Device, it is assumed that Devices host only one or very few device applications, which most of the time communicate each to a single network application at a time.

The radio protocol authenticates and ensures the integrity of each message. This is achieved by using a unique device ID and an AES-128 based message authentication code, ensuring that the message has been generated and sent by the device with the ID claimed in the message. Application data can be encrypted at the application level or not, depending on the criticality of the use case, to provide a balance between cost and effort vs. risk. AES-128 in counter mode is used for encryption. Cryptographic keys are independent for each device. These keys are associated with the device ID and separate integrity and confidentiality keys are pre-provisioned. A confidentiality key is only provisioned if confidentiality is to be used. At the time of writing the algorithms and keying details for this are not published.

#### 2.4. Wi-SUN Alliance Field Area Network (FAN)

Text here is via personal communication from Bob Heile (bheile@ieee.org) and was authored by Bob and Sum Chin Sean. Duffy (paduffy@cisco.com) also provided additional comments/input on this section.

##### 2.4.1. Provenance and Documents

The Wi-SUN Alliance <<https://www.wi-sun.org/>> is an industry alliance for smart city, smart grid, smart utility, and a broad set of general IoT applications. The Wi-SUN Alliance Field Area Network (FAN) profile is open standards based (primarily on IETF and IEEE802 standards) and was developed to address applications like smart municipality/city infrastructure monitoring and management, electric vehicle (EV) infrastructure, advanced metering infrastructure (AMI), distribution automation (DA), supervisory control and data acquisition (SCADA) protection/management, distributed generation monitoring and management, and many more IoT applications. Additionally, the Alliance has created a certification program to promote global multi-vendor interoperability.

The FAN profile is specified within ANSI/TIA as an extension of work previously done on Smart Utility Networks. [ANSI-4957-000]. Updates to those specifications intended to be published in 2017 will contain details of the FAN profile. A current snapshot of the work to

produce that profile is presented in [wisun-pressie1]  
[wisun-pressie2] .

#### 2.4.2. Characteristics

The FAN profile is an IPv6 wireless mesh network with support for enterprise level security. The frequency hopping wireless mesh topology aims to offer superior network robustness, reliability due to high redundancy, good scalability due to the flexible mesh configuration and good resilience to interference. Very low power modes are in development permitting long term battery operation of network nodes.

The following list contains some overall characteristics of Wi-SUN that are relevant to LPWAN applications.

- o Coverage: The range of Wi-SUN FAN is typically 2 -- 3 km in line of sight, matching the needs of neighborhood area networks, campus area networks, or corporate area networks. The range can also be extended via multi-hop networking.
- o High bandwidth, low link latency: Wi-SUN supports relatively high bandwidth, i.e. up to 300 kbps [FANTPS], enables remote update and upgrade of devices so that they can handle new applications, extending their working life. Wi-SUN supports LPWAN IoT applications that require on-demand control by providing low link latency (0.02s) and bi-directional communication.
- o Low power consumption: FAN devices draw less than 2 uA when resting and only 8 mA when listening. Such devices can maintain a long lifetime even if they are frequently listening. For instance, suppose the device transmits data for 10 ms once every 10 s; theoretically, a battery of 1000 mAh can last more than 10 years.
- o Scalability: Tens of millions Wi-SUN FAN devices have been deployed in urban, suburban and rural environments, including deployments with more than 1 million devices.

A FAN contains one or more networks. Within a network, nodes assume one of three operational roles. First, each network contains a Border Router providing Wide Area Network (WAN) connectivity to the network. The Border Router maintains source routing tables for all nodes within its network, provides node authentication and key management services, and disseminates network-wide information such as broadcast schedules. Secondly, Router nodes, which provide upward and downward packet forwarding (within a network). A Router also provides services for relaying security and address management

protocols. Lastly, Leaf nodes provide minimum capabilities: discovering and joining a network, send/receive IPv6 packets, etc. A low power network may contain a mesh topology with Routers at the edges that construct a star topology with Leaf nodes.

The FAN profile is based on various open standards developed by the IETF (including [RFC0768], [RFC2460], [RFC4443] and [RFC6282]), IEEE802 (including [IEEE-802-15-4] and [IEEE-802-15-9]) and ANSI/TIA [ANSI-4957-210] for low power and lossy networks.

The FAN profile specification provides an application-independent IPv6-based transport service. There are two possible methods for establishing the IPv6 packet routing: Routing Protocol for Low-Power and Lossy Networks (RPL) at the Network layer is mandatory, and Multi-Hop Delivery Service (MHDS) is optional at the Data Link layer. Table 5 provides an overview of the FAN network stack.

The Transport service is based on User Datagram Protocol (UDP) defined in RFC768 or Transmission Control Protocol (TCP) defined in RFC793.

The Network service is provided by IPv6 as defined in RFC2460 with 6LoWPAN adaptation as defined in RFC4944 and RFC6282. ICMPv6, as defined in RFC4443, is used for the control plane during information exchange.

The Data Link service provides both control/management of the Physical layer and data transfer/management services to the Network layer. These services are divided into Media Access Control (MAC) and Logical Link Control (LLC) sub-layers. The LLC sub-layer provides a protocol dispatch service which supports 6LoWPAN and an optional MAC sub-layer mesh service. The MAC sub-layer is constructed using data structures defined in IEEE802.15.4-2015. Multiple modes of frequency hopping are defined. The entire MAC payload is encapsulated in an IEEE802.15.9 Information Element to enable LLC protocol dispatch between upper layer 6LoWPAN processing, MAC sublayer mesh processing, etc. These areas will be expanded once IEEE802.15.12 is completed.

The PHY service is derived from a sub-set of the SUN FSK specification in IEEE802.15.4-2015. The 2-FSK modulation schemes, with channel spacing range from 200 to 600 kHz, are defined to provide data rates from 50 to 300 kbps, with Forward Error Coding (FEC) as an optional feature. Towards enabling ultra-low-power applications, the PHY layer design is also extendable to low energy and critical infrastructure monitoring networks.

Layer	Description
IPv6 protocol suite	TCP/UDP 6LoWPAN Adaptation + Header Compression DHCPv6 for IP address management. Routing using RPL. ICMPv6. Unicast and Multicast forwarding.
MAC based on IEEE 802.15.4e + IE extensions	Frequency hopping Discovery and Join Protocol Dispatch (IEEE 802.15.9) Several Frame Exchange patterns Optional Mesh Under routing (ANSI 4957.210).
PHY based on 802.15.4g	Various data rates and regions
Security	802.1X/EAP-TLS/PKI Authentication. TLS_ECDHE_ECDSA_WITH_AES_128_CCM_8 required for EAP-TLS. 802.11i Group Key Management Frame security is implemented as AES-CCM* as specified in IEEE 802.15.4 Optional ETSI-TS-102-887-2 Node 2 Node Key Management

Table 5: Wi-SUN Stack Overview

The FAN security supports Data Link layer network access control, mutual authentication, and establishment of a secure pairwise link

between a FAN node and its Border Router, which is implemented with an adaptation of IEEE802.1X and EAP-TLS as described in [RFC5216] using secure device identity as described in IEEE802.1AR. Certificate formats are based upon [RFC5280]. A secure group link between a Border Router and a set of FAN nodes is established using an adaptation of the IEEE802.11 Four-Way Handshake. A set of 4 group keys are maintained within the network, one of which is the current transmit key. Secure node to node links are supported between one-hop FAN neighbors using an adaptation of ETSI-TS-102-887-2. FAN nodes implement Frame Security as specified in IEEE802.15.4-2015.

### 3. Generic Terminology

LPWAN technologies, such as those discussed above, have similar architectures but different terminology. We can identify different types of entities in a typical LPWAN network:

- o End-Devices are the devices or the "things" (e.g. sensors, actuators, etc.); they are named differently in each technology (End Device, User Equipment or End Point). There can be a high density of end devices per radio gateway.
- o The Radio Gateway, which is the end point of the constrained link. It is known as: Gateway, Evolved Node B or Base station.
- o The Network Gateway or Router is the interconnection node between the Radio Gateway and the Internet. It is known as: Network Server, Serving GW or Service Center.
- o LPWAN-AAA Server, which controls the user authentication, the applications. It is known as: Join-Server, Home Subscriber Server or Registration Authority. (We use the term LPWAN-AAA server because we're not assuming that this entity speaks RADIUS or Diameter as many/most AAA servers do, but equally we don't want to rule that out, as the functionality will be similar.
- o At last we have the Application Server, known also as Packet Data Node Gateway or Network Application.



Function/ Technology	LORAWAN	NB-IOT	SIGFOX	Wi-SUN	IETF
Sensor, Actuator, device, object	End Device	User Equipment	End Point	Leaf Node	Device (Dev)
Transceiver Antenna	Gateway	Evolved Node B	Base Station	Router Node	RADIO Gateway
Server	Network Server	PDN GW/ SCEF	Service Center	Border Router	Network Gateway (NGW)
Security Server	Join Server	Home Subscriber Server	Registration Authority	Authent. Server	LPWAN- AAA SERVER
Application	Application Server	Application Server	Network Application	Appli- cation	Application (App)

Figure 8: LPWAN Architecture Terminology




Figure 9: LPWAN Architecture

In addition to the names of entities, LPWANs are also subject to possibly regional frequency band regulations. Those may include restrictions on the duty-cycle, for example requiring that hosts only transmit for a certain percentage of each hour.

#### 4. Gap Analysis

This section considers some of the gaps between current LPWAN technologies and the goals of the LPWAN working group. Many of the generic considerations described in [RFC7452] will also apply in LPWANs, as end-devices can also be considered as a subclass of (so-

called) "smart objects." In addition, LPWAN device implementers will also need to consider the issues relating to firmware updates described in [RFC8240].

#### 4.1. Naive application of IPv6

IPv6 [RFC8200] has been designed to allocate addresses to all the nodes connected to the Internet. Nevertheless, the header overhead of at least 40 bytes introduced by the protocol is incompatible with LPWAN constraints. If IPv6 with no further optimization were used, several LPWAN frames could be needed just to carry the IP header. Another problem arises from IPv6 MTU requirements, which require the layer below to support at least 1280 byte packets [RFC2460].

IPv6 has a configuration protocol - neighbor discovery protocol, (NDP) [RFC4861]). For a node to learn network parameters NDP generates regular traffic with a relatively large message size that does not fit LPWAN constraints.

In some LPWAN technologies, layer two multicast is not supported. In that case, if the network topology is a star, the solution and considerations of section 3.2.5 of [RFC7668] may be applied.

Other key protocols such as DHCPv6 [RFC3315], IPsec [RFC4301] and TLS [RFC5246] have similarly problematic properties in this context. Each of those require relatively frequent round-trips between the host and some other host on the network. In the case of cryptographic protocols such as IPsec and TLS, in addition to the round-trips required for secure session establishment, cryptographic operations can require padding and addition of authenticators that are problematic when considering LPWAN lower layers. Note that mains powered Wi-SUN mesh router nodes will typically be more resource capable than the other LPWAN techs discussed. This can enable use of more "chatty" protocols for some aspects of Wi-SUN.

#### 4.2. 6LoWPAN

Several technologies that exhibit significant constraints in various dimensions have exploited the 6LoWPAN suite of specifications [RFC4944], [RFC6282], [RFC6775] to support IPv6 [I-D.hong-6lo-use-cases]. However, the constraints of LPWANs, often more extreme than those typical of technologies that have (re)used 6LoWPAN, constitute a challenge for the 6LoWPAN suite in order to enable IPv6 over LPWAN. LPWANs are characterized by device constraints (in terms of processing capacity, memory, and energy availability), and specially, link constraints, such as:

- o tiny layer two payload size (from ~10 to ~100 bytes),

- o very low bit rate (from ~10 bit/s to ~100 kbit/s), and
- o in some specific technologies, further message rate constraints (e.g. between ~0.1 message/minute and ~1 message/minute) due to regional regulations that limit the duty cycle.

#### 4.2.1. Header Compression

6LoWPAN header compression reduces IPv6 (and UDP) header overhead by eliding header fields when they can be derived from the link layer, and by assuming that some of the header fields will frequently carry expected values. 6LoWPAN provides both stateless and stateful header compression. In the latter, all nodes of a 6LoWPAN are assumed to share compression context. In the best case, the IPv6 header for link-local communication can be reduced to only 2 bytes. For global communication, the IPv6 header may be compressed down to 3 bytes in the most extreme case. However, in more practical situations, the smallest IPv6 header size may be 11 bytes (one address prefix compressed) or 19 bytes (both source and destination prefixes compressed). These headers are large considering the link layer payload size of LPWAN technologies, and in some cases are even bigger than the LPWAN PDUs. 6LoWPAN has been initially designed for IEEE 802.15.4 networks with a frame size up to 127 bytes and a throughput of up to 250 kb/s, which may or may not be duty-cycled.

#### 4.2.2. Address Autoconfiguration

Traditionally, Interface Identifiers (IIDs) have been derived from link layer identifiers [RFC4944]. This allows optimizations such as header compression. Nevertheless, recent guidance has given advice on the fact that, due to privacy concerns, 6LoWPAN devices should not be configured to embed their link layer addresses in the IID by default. [RFC8065] provides guidance on better methods for generating IIDs.

#### 4.2.3. Fragmentation

As stated above, IPv6 requires the layer below to support an MTU of 1280 bytes [RFC2460]. Therefore, given the low maximum payload size of LPWAN technologies, fragmentation is needed.

If a layer of an LPWAN technology supports fragmentation, proper analysis has to be carried out to decide whether the fragmentation functionality provided by the lower layer or fragmentation at the adaptation layer should be used. Otherwise, fragmentation functionality shall be used at the adaptation layer.

6LoWPAN defined a fragmentation mechanism and a fragmentation header to support the transmission of IPv6 packets over IEEE 802.15.4 networks [RFC4944]. While the 6LoWPAN fragmentation header is appropriate for IEEE 802.15.4-2003 (which has a frame payload size of 81-102 bytes), it is not suitable for several LPWAN technologies, many of which have a maximum payload size that is one order of magnitude below that of IEEE 802.15.4-2003. The overhead of the 6LoWPAN fragmentation header is high, considering the reduced payload size of LPWAN technologies and the limited energy availability of the devices using such technologies. Furthermore, its datagram offset field is expressed in increments of eight octets. In some LPWAN technologies, the 6LoWPAN fragmentation header plus eight octets from the original datagram exceeds the available space in the layer two payload. In addition, the MTU in the LPWAN networks could be variable which implies a variable fragmentation solution.

#### 4.2.4. Neighbor Discovery

6LoWPAN Neighbor Discovery [RFC6775] defined optimizations to IPv6 Neighbor Discovery [RFC4861], in order to adapt functionality of the latter for networks of devices using IEEE 802.15.4 or similar technologies. The optimizations comprise host-initiated interactions to allow for sleeping hosts, replacement of multicast-based address resolution for hosts by an address registration mechanism, multihop extensions for prefix distribution and duplicate address detection (note that these are not needed in a star topology network), and support for 6LoWPAN header compression.

6LoWPAN Neighbor Discovery may be used in not so severely constrained LPWAN networks. The relative overhead incurred will depend on the LPWAN technology used (and on its configuration, if appropriate). In certain LPWAN setups (with a maximum payload size above ~60 bytes, and duty-cycle-free or equivalent operation), an RS/RA/NS/NA exchange may be completed in a few seconds, without incurring packet fragmentation.

In other LPWANs (with a maximum payload size of ~10 bytes, and a message rate of ~0.1 message/minute), the same exchange may take hours or even days, leading to severe fragmentation and consuming a significant amount of the available network resources. 6LoWPAN Neighbor Discovery behavior may be tuned through the use of appropriate values for the default Router Lifetime, the Valid Lifetime in the PIOs, and the Valid Lifetime in the 6LoWPAN Context Option (6CO), as well as the address Registration Lifetime. However, for the latter LPWANs mentioned above, 6LoWPAN Neighbor Discovery is not suitable.

#### 4.3. 6lo

The 6lo WG has been reusing and adapting 6LoWPAN to enable IPv6 support over link layer technologies such as Bluetooth Low Energy (BTLE), ITU-T G.9959, DECT-ULE, MS/TP-RS485, NFC IEEE 802.11ah. (See <<https://tools.ietf.org/wg/6lo>> for details.) These technologies are similar in several aspects to IEEE 802.15.4, which was the original 6LoWPAN target technology.

6lo has mostly used the subset of 6LoWPAN techniques best suited for each lower layer technology, and has provided additional optimizations for technologies where the star topology is used, such as BTLE or DECT-ULE.

The main constraint in these networks comes from the nature of the devices (constrained devices), whereas in LPWANs it is the network itself that imposes the most stringent constraints.

#### 4.4. 6tisch

The 6tisch solution is dedicated to mesh networks that operate using 802.15.4e MAC with a deterministic slotted channel. The time slot channel (TSCH) can help to reduce collisions and to enable a better balance over the channels. It improves the battery life by avoiding the idle listening time for the return channel.

A key element of 6tisch is the use of synchronization to enable determinism. TSCH and 6TiSCH may provide a standard scheduling function. The LPWAN networks probably will not support synchronization like the one used in 6tisch.

#### 4.5. RoHC

Robust header compression (RoHC) is a header compression mechanism [RFC3095] developed for multimedia flows in a point to point channel. RoHC uses 3 levels of compression, each level having its own header format. In the first level, RoHC sends 52 bytes of header, in the second level the header could be from 34 to 15 bytes and in the third level header size could be from 7 to 2 bytes. The level of compression is managed by a sequence number, which varies in size from 2 bytes to 4 bits in the minimal compression. SN compression is done with an algorithm called W-LSB (Window- Least Significant Bits). This window has a 4-bit size representing 15 packets, so every 15 packets RoHC needs to slide the window in order to receive the correct sequence number, and sliding the window implies a reduction of the level of compression. When packets are lost or errored, the decompressor loses context and drops packets until a bigger header is sent with more complete information. To estimate the performance of

RoHC, an average header size is used. This average depends on the transmission conditions, but most of the time is between 3 and 4 bytes.

RoHC has not been adapted specifically to the constrained hosts and networks of LPWANs: it does not take into account energy limitations nor the transmission rate, and RoHC context is synchronised during transmission, which does not allow better compression.

#### 4.6. ROLL

Most technologies considered by the lpwan WG are based on a star topology, which eliminates the need for routing at that layer. Future work may address additional use-cases that may require adaptation of existing routing protocols or the definition of new ones. As of the time of writing, work similar to that done in the ROLL WG and other routing protocols are out of scope of the LPWAN WG.

#### 4.7. CoAP

CoAP [RFC7252] provides a RESTful framework for applications intended to run on constrained IP networks. It may be necessary to adapt CoAP or related protocols to take into account for the extreme duty cycles and the potentially extremely limited throughput of LPWANs.

For example, some of the timers in CoAP may need to be redefined. Taking into account CoAP acknowledgments may allow the reduction of L2 acknowledgments. On the other hand, the current work in progress in the CoRE WG where the COMI/CoOL network management interface which, uses Structured Identifiers (SID) to reduce payload size over CoAP may prove to be a good solution for the LPWAN technologies. The overhead is reduced by adding a dictionary which matches a URI to a small identifier and a compact mapping of the YANG model into the CBOR binary representation.

#### 4.8. Mobility

LPWAN nodes can be mobile. However, LPWAN mobility is different from the one specified for Mobile IP. LPWAN implies sporadic traffic and will rarely be used for high-frequency, real-time communications. The applications do not generate a flow, they need to save energy and most of the time the node will be down.

In addition, LPWAN mobility may mostly apply to groups of devices, that represent a network in which case mobility is more a concern for the gateway than the devices. NEMO [RFC3963] Mobility or other mobile gateway solutions (such as a gateway with an LTE uplink) may be used in the case where some end-devices belonging to the same

network gateway move from one point to another such that they are not aware of being mobile.

#### 4.9. DNS and LPWAN

The Domain Name System (DNS) [RFC1035], enables applications to name things with a globally resolvable name. Many protocols use the DNS to identify hosts, for example applications using CoAP.

The DNS query/answer protocol as a pre-cursor to other communication within the time-to-live (TTL) of a DNS answer is clearly problematic in an LPWAN, say where only one round-trip per hour can be used, and with a TTL that is less than 3600. It is currently unclear whether and how DNS-like functionality might be provided in LPWANs.

#### 5. Security Considerations

Most LPWAN technologies integrate some authentication or encryption mechanisms that were defined outside the IETF. The working group may need to do work to integrate these mechanisms to unify management. A standardized Authentication, Accounting, and Authorization (AAA) infrastructure [RFC2904] may offer a scalable solution for some of the security and management issues for LPWANs. AAA offers centralized management that may be of use in LPWANs, for example [I-D.garcia-dime-diameter-lorawan] and [I-D.garcia-radext-radius-lorawan] suggest possible security processes for a LoRaWAN network. Similar mechanisms may be useful to explore for other LPWAN technologies.

Some applications using LPWANs may raise few or no privacy considerations. For example, temperature sensors in a large office building may not raise privacy issues. However, the same sensors, if deployed in a home environment and especially if triggered due to human presence, can raise significant privacy issues - if an end-device emits (an encrypted) packet every time someone enters a room in a home, then that traffic is privacy sensitive. And the more that the existence of that traffic is visible to network entities, the more privacy sensitivities arise. At this point, it is not clear whether there are workable mitigations for problems like this - in a more typical network, one would consider defining padding mechanisms and allowing for cover traffic. In some LPWANs, those mechanisms may not be feasible. Nonetheless, the privacy challenges do exist and can be real and so some solutions will be needed. Note that many aspects of solutions in this space may not be visible in IETF specifications, but can be e.g. implementation or deployment specific.

Another challenge for LPWANs will be how to handle key management and associated protocols. In a more traditional network (e.g. the web), servers can "staple" Online Certificate Status Protocol (OCSP) responses in order to allow browsers to check revocation status for presented certificates. [RFC6961] While the stapling approach is likely something that would help in an LPWAN, as it avoids an RTT, certificates and OCSP responses are bulky items and will prove challenging to handle in LPWANs with bounded bandwidth.

## 6. IANA Considerations

There are no IANA considerations related to this memo.

## 7. Contributors

[[RFC editor: Please fix names below for I18N.]]

As stated above this document is mainly a collection of content developed by the full set of contributors listed below. The main input documents and their authors were:

- o Text for Section 2.1 was provided by Alper Yegin and Stephen Farrell in [I-D.farrell-lpwan-lora-overview].
- o Text for Section 2.2 was provided by Antti Ratilainen in [I-D.ratilainen-lpwan-nb-iot].
- o Text for Section 2.3 was provided by Juan Carlos Zuniga and Benoit Ponsard in [I-D.zuniga-lpwan-sigfox-system-description].
- o Text for Section 2.4 was provided via personal communication from Bob Heile (bheile@ieee.org) and was authored by Bob and Sum Chin Sean. There is no Internet draft for that at present.
- o Text for Section 4 was provided by Ana Minabiru, Carles Gomez, Laurent Toutain, Josep Paradells and Jon Crowcroft in [I-D.minaburo-lpwan-gap-analysis]. Additional text from that draft is also used elsewhere above.

The full list of contributors are:

Jon Crowcroft  
University of Cambridge  
JJ Thomson Avenue  
Cambridge, CB3 0FD  
United Kingdom



Email: [jon.crowcroft@cl.cam.ac.uk](mailto:jon.crowcroft@cl.cam.ac.uk)

Carles Gomez  
UPC/i2CAT  
C/Esteve Terradas, 7  
Castelldefels 08860  
Spain

Email: [carlesgo@entel.upc.edu](mailto:carlesgo@entel.upc.edu)

Bob Heile  
Wi-Sun Alliance  
11 Robert Toner Blvd, Suite 5-301  
North Attleboro, MA 02763  
USA

Phone: +1-781-929-4832  
Email: [bheile@ieee.org](mailto:bheile@ieee.org)

Ana Minaburo  
Acklio  
2bis rue de la Chataigneraie  
35510 Cesson-Sevigne Cedex  
France

Email: [ana@ackl.io](mailto:ana@ackl.io)

Josep PARadells  
UPC/i2CAT  
C/Jordi Girona, 1-3  
Barcelona 08034  
Spain

Email: [josep.paradells@entel.upc.edu](mailto:josep.paradells@entel.upc.edu)

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

Email: [charliep@computer.org](mailto:charliep@computer.org)

Benoit Ponsard  
SIGFOX  
425 rue Jean Rostand  
Labege 31670  
France

Email: [Benoit.Ponsard@sigfox.com](mailto:Benoit.Ponsard@sigfox.com)  
URI: <http://www.sigfox.com/>

Antti Ratilainen  
Ericsson  
Hirsalantie 11  
Jorvas 02420  
Finland

Email: [antti.ratilainen@ericsson.com](mailto:antti.ratilainen@ericsson.com)

Chin-Sean SUM  
Wi-Sun Alliance  
20, Science Park Rd  
Singapore 117674

Phone: +65 6771 1011  
Email: [sum@wi-sun.org](mailto:sum@wi-sun.org)

Laurent Toutain  
Institut MINES TELECOM ; TELECOM Bretagne  
2 rue de la Chataigneraie  
CS 17607  
35576 Cesson-Sevigne Cedex  
France

Email: [Laurent.Toutain@telecom-bretagne.eu](mailto:Laurent.Toutain@telecom-bretagne.eu)

Alper Yegin  
Actility  
Paris, Paris  
FR

Email: [alper.yegin@actility.com](mailto:alper.yegin@actility.com)

Juan Carlos Zuniga  
SIGFOX

425 rue Jean Rostand  
Labege 31670  
France

Email: JuanCarlos.Zuniga@sigfox.com  
URI: <http://www.sigfox.com/>

## 8. Acknowledgments

Thanks to all those listed in Section 7 for the excellent text. Errors in the handling of that are solely the editor's fault.

[[RFC editor: Please fix names below for I18N, at least Mirja's does need fixing.]]

In addition to the contributors above, thanks are due to (in alphabetical order): Abdussalam Baryun, Andy Malis, Arun (arun@acklio.com), Behcet SariKaya, Dan Garcia Carrillo, Jiazi Yi, Mirja Kuehlewind, Paul Duffy, Russ Housley, Samita Chakrabarti, Thad Guidry, Warren Kumari, for comments.

Alexander Pelov and Pascal Thubert were the LPWAN WG chairs while this document was developed.

Stephen Farrell's work on this memo was supported by Pervasive Nation, the Science Foundation Ireland's CONNECT centre national IoT network. <<https://connectcentre.ie/pervasive-nation/>>

## 9. Informative References

- [RFC0768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, DOI 10.17487/RFC0768, August 1980, <<https://www.rfc-editor.org/info/rfc768>>.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <<https://www.rfc-editor.org/info/rfc1035>>.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <<https://www.rfc-editor.org/info/rfc2460>>.
- [RFC2904] Vollbrecht, J., Calhoun, P., Farrell, S., Gommans, L., Gross, G., de Bruijn, B., de Laat, C., Holdrege, M., and D. Spence, "AAA Authorization Framework", RFC 2904, DOI 10.17487/RFC2904, August 2000, <<https://www.rfc-editor.org/info/rfc2904>>.

- [RFC3095] Bormann, C., Burmeister, C., Degermark, M., Fukushima, H., Hannu, H., Jonsson, L-E., Hakenberg, R., Koren, T., Le, K., Liu, Z., Martensson, A., Miyazaki, A., Svanbro, K., Wiebke, T., Yoshimura, T., and H. Zheng, "RObust Header Compression (ROHC): Framework and four profiles: RTP, UDP, ESP, and uncompressed", RFC 3095, DOI 10.17487/RFC3095, July 2001, <<https://www.rfc-editor.org/info/rfc3095>>.
- [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>>.
- [RFC3963] Devarapalli, V., Wakikawa, R., Petrescu, A., and P. Thubert, "Network Mobility (NEMO) Basic Support Protocol", RFC 3963, DOI 10.17487/RFC3963, January 2005, <<https://www.rfc-editor.org/info/rfc3963>>.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <<https://www.rfc-editor.org/info/rfc4301>>.
- [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>>.
- [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>>.
- [RFC5216] Simon, D., Aboba, B., and R. Hurst, "The EAP-TLS Authentication Protocol", RFC 5216, DOI 10.17487/RFC5216, March 2008, <<https://www.rfc-editor.org/info/rfc5216>>.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <<https://www.rfc-editor.org/info/rfc5246>>.

- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <<https://www.rfc-editor.org/info/rfc5280>>.
- [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>>.
- [RFC6961] Pettersen, Y., "The Transport Layer Security (TLS) Multiple Certificate Status Request Extension", RFC 6961, DOI 10.17487/RFC6961, June 2013, <<https://www.rfc-editor.org/info/rfc6961>>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/RFC7252, June 2014, <<https://www.rfc-editor.org/info/rfc7252>>.
- [RFC7452] Tschofenig, H., Arkko, J., Thaler, D., and D. McPherson, "Architectural Considerations in Smart Object Networking", RFC 7452, DOI 10.17487/RFC7452, March 2015, <<https://www.rfc-editor.org/info/rfc7452>>.
- [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>>.
- [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>>.
- [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>>.

- [RFC8240]    Tschofenig, H. and S. Farrell, "Report from the Internet of Things Software Update (IoTSU) Workshop 2016", RFC 8240, DOI 10.17487/RFC8240, September 2017, <<https://www.rfc-editor.org/info/rfc8240>>.
- [I-D.farrell-lpwan-lora-overview]  
Farrell, S. and A. Yegin, "LoRaWAN Overview", draft-farrell-lpwan-lora-overview-01 (work in progress), October 2016.
- [I-D.minaburo-lpwan-gap-analysis]  
Minaburo, A., Gomez, C., Toutain, L., Paradells, J., and J. Crowcroft, "LPWAN Survey and GAP Analysis", draft-minaburo-lpwan-gap-analysis-02 (work in progress), October 2016.
- [I-D.zuniga-lpwan-sigfox-system-description]  
Zuniga, J. and B. PONSARD, "SIGFOX System Description", draft-zuniga-lpwan-sigfox-system-description-04 (work in progress), December 2017.
- [I-D.ratilainen-lpwan-nb-iot]  
Ratilainen, A., "NB-IoT characteristics", draft-ratilainen-lpwan-nb-iot-00 (work in progress), July 2016.
- [I-D.garcia-dime-diameter-lorawan]  
Garcia, D., Lopez, R., Kandasamy, A., and A. Pelov, "LoRaWAN Authentication in Diameter", draft-garcia-dime-diameter-lorawan-00 (work in progress), May 2016.
- [I-D.garcia-radext-radius-lorawan]  
Garcia, D., Lopez, R., Kandasamy, A., and A. Pelov, "LoRaWAN Authentication in RADIUS", draft-garcia-radext-radius-lorawan-03 (work in progress), May 2017.
- [I-D.hong-6lo-use-cases]  
Hong, Y. and C. Gomez, "IPv6 over Constrained Node Networks(6lo) Applicability & Use cases", draft-hong-6lo-use-cases-03 (work in progress), October 2016.
- [TGPP36300]  
3GPP, "TS 36.300 v13.4.0 Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Overall description; Stage 2", 2016, <[http://www.3gpp.org/ftp/Specs/2016-09/Rel-14/36\\_series/](http://www.3gpp.org/ftp/Specs/2016-09/Rel-14/36_series/)>.

- [TGPP36321]  
3GPP, "TS 36.321 v13.2.0 Evolved Universal Terrestrial Radio Access (E-UTRA); Medium Access Control (MAC) protocol specification", 2016.
- [TGPP36322]  
3GPP, "TS 36.322 v13.2.0 Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Link Control (RLC) protocol specification", 2016.
- [TGPP36323]  
3GPP, "TS 36.323 v13.2.0 Evolved Universal Terrestrial Radio Access (E-UTRA); Packet Data Convergence Protocol (PDCP) specification (Not yet available)", 2016.
- [TGPP36331]  
3GPP, "TS 36.331 v13.2.0 Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Resource Control (RRC); Protocol specification", 2016.
- [TGPP36201]  
3GPP, "TS 36.201 v13.2.0 - Evolved Universal Terrestrial Radio Access (E-UTRA); LTE physical layer; General description", 2016.
- [TGPP23720]  
3GPP, "TR 23.720 v13.0.0 - Study on architecture enhancements for Cellular Internet of Things", 2016.
- [TGPP33203]  
3GPP, "TS 33.203 v13.1.0 - 3G security; Access security for IP-based services", 2016.
- [fcc\_ref] "FCC CFR 47 Part 15.247 Telecommunication Radio Frequency Devices - Operation within the bands 902-928 MHz, 2400-2483.5 MHz, and 5725-5850 MHz.", June 2016.
- [etsi\_ref]  
"ETSI EN 300-220 (Parts 1 and 2): Electromagnetic compatibility and Radio spectrum Matters (ERM); Short Range Devices (SRD); Radio equipment to be used in the 25 MHz to 1 000 MHz frequency range with power levels ranging up to 500 mW", May 2016.
- [arib\_ref]  
"ARIB STD-T108 (Version 1.0): 920MHz-Band Telemeter, Telecontrol and data transmission radio equipment.", February 2012.

[LoRaSpec]

LoRa Alliance, "LoRaWAN Specification Version V1.0.2", July 2016, <[http://portal.lora-alliance.org/DesktopModules/Inventures\\_Document/FileDownload.aspx?ContentID=1398](http://portal.lora-alliance.org/DesktopModules/Inventures_Document/FileDownload.aspx?ContentID=1398)>.

[ANSI-4957-000]

ANSI, TIA-4957.000, "Architecture Overview for the Smart Utility Network", May 2013, <[https://global.ihs.com/doc\\_detail.cfm?%26rid=TIA%26item\\_s\\_key=00606368](https://global.ihs.com/doc_detail.cfm?%26rid=TIA%26item_s_key=00606368)>.

[ANSI-4957-210]

ANSI, TIA-4957.210, "Multi-Hop Delivery Specification of a Data Link Sub-Layer", May 2013, <[https://global.ihs.com/doc\\_detail.cfm?%26csf=TIA%26item\\_s\\_key=00601800](https://global.ihs.com/doc_detail.cfm?%26csf=TIA%26item_s_key=00601800)>.

[wisun-pressie1]

Phil Beecher, Chair, Wi-SUN Alliance, "Wi-SUN Alliance Overview", March 2017, <<http://indiasmartgrid.org/event2017/10-03-2017/4.%20Roundtable%20on%20Communication%20and%20Cyber%20Security/1.%20Phil%20Beecher.pdf>>.

[wisun-pressie2]

Bob Heile, Director of Standards, Wi-SUN Alliance, "IETF97 Wi-SUN Alliance Field Area Network (FAN) Overview", November 2016, <<https://www.ietf.org/proceedings/97/slides/slides-97-lpwan-35-wi-sun-presentation-00.pdf>>.

[IEEE-802-15-4]

"IEEE Standard for Low-Rate Wireless Personal Area Networks (WPANs)", IEEE Standard 802.15.4, 2015, <<https://standards.ieee.org/findstds/standard/802.15.4-2015.html>>.

[IEEE-802-15-9]

"IEEE Recommended Practice for Transport of Key Management Protocol (KMP) Datagrams", IEEE Standard 802.15.9, 2016, <<https://standards.ieee.org/findstds/standard/802.15.9-2016.html>>.

[etsi\_unb]

"ETSI TR 103 435 System Reference document (SRdoc); Short Range Devices (SRD); Technical characteristics for Ultra Narrow Band (UNB) SRDs operating in the UHF spectrum below 1 GHz", February 2017.



[nbiot-ov]

Beyene, Yihene Dagne, et al., "NB-IoT technology overview and experience from cloud-RAN implementation", IEEE Wireless Communications 24,3 (2017): 26-32, June 2017.

## Appendix A. Changes

[[RFC editor: Please remove this before publication]]

### A.1. From -00 to -01

- o WG have stated they want this to be an RFC.
- o WG clearly want to keep the RF details.
- o Various changes made to remove/resolve a number of editorial notes from -00 (in some cases as per suggestions from Ana Minaburo)
- o Merged PR's: #1...
- o Rejected PR's: #2 (change was made to .txt not .xml but was replicated manually by editor)
- o Github repo is at: <https://github.com/sftcd/lpwan-ov>

### A.2. From -01 to -02

- o WG seem to agree with editor suggestions in slides 13-24 of the presentation on this topic given at IETF98 (See: <https://www.ietf.org/proceedings/98/slides/slides-98-lpwan-aggregated-slides-07.pdf>)
- o Got new text wrt Wi-SUN via email from Paul Duffy and merged that in
- o Reflected list discussion wrt terminology and "end-device"
- o Merged PR's: #3...

### A.3. From -02 to -03

- o Editorial changes and typo fixes thanks to Fred Baker running something called Grammerly and sending me it's report.
- o Merged PR's: #4, #6, #7...
- o Editor did an editing pass on the lot.

A.4. From -03 to -04

- o Picked up a PR that had been wrongly applied that expands UE
- o Editorial changes wrt LoRa suggested by Alper
- o Editorial changes wrt SIGFOX provided by Juan-Carlos

A.5. From -04 to -05

- o Handled Russ Housley's WGLC review.
- o Handled Alper Yegin's WGLC review.

A.6. From -05 to -06

- o More Alper comments:-)
- o Added some more detail about sigfox security.
- o Added Wi-SUN changes from Charlie Perkins

A.7. From -06 to -07

Yet more Alper comments:-)

Comments from Behcet Sarikaya

A.8. From -07 to -08

various typos

Last call and directorate comments from Abdussalam Baryun (AB) and Andy Malis

20180118 IESG ballot comments from Warren: nits handled, two possible bits of text still needed.

Some more AB comments handled. Still need to check over 7452 and 8240 to see if issues from those need to be discussed here.

Corrected "no IP capabilities - Wi-SUN devices do v6 (thanks Paul Duffy:-)

Mirja's AD ballot comments handled.

Added a sentence in intro trying to say what's "special" about LPWAN compared to other constrained networks. (As suggested by Warren.)

Added text @ start of gap analysis referring to RFCs 7252 and 8240, as suggested by a few folks (AB, Warren, Mirja)

Added nbiot-ov reference for those who'd like a more polished presentation of NB-IoT

A.9. From -08 to -09

Changes due to IoT-DIR review from Samita Chakrabarti: fixed error on max rate between tables 1 and 2; s/eNb/eNodeB/; fixed references to hong-6lo-use-cases; added RFC8065 reference

A.10. From -09 to -10

Added Charlie Perkins as contributor - was supposed to have been done ages ago - editor forgot;-)

Author's Address

Stephen Farrell (editor)  
Trinity College Dublin  
Dublin 2  
Ireland

Phone: +353-1-896-2354  
Email: stephen.farrell@cs.tcd.ie

lpwan Working Group  
Internet-Draft  
Intended status: Informational  
Expires: January 3, 2019

N. Sornin, Ed.  
M. Coracin  
Semtech  
I. Petrov  
Acklio  
A. Yegin  
Actility  
J. Catalano  
Kerlink  
V. Audebert  
EDF R&D  
July 02, 2018

Static Context Header Compression (SCHC) over LoRaWAN  
draft-petrov-lpwan-ipv6-schc-over-lorawan-02

Abstract

The Static Context Header Compression (SCHC) specification describes generic header compression and fragmentation techniques for LPWAN (Low Power Wide Area Networks) technologies. SCHC is a generic mechanism designed for great flexibility, so that it can be adapted for any of the LPWAN technologies.

This document provides the adaptation of SCHC for use in LoRaWAN networks, and provides elements such as efficient parameterization and modes of operation.

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
2. Terminology . . . . .	3
3. Static Context Header Compression Overview . . . . .	3
4. LoRaWAN Architecture . . . . .	4
4.1. Device classes (A, B, C) and interactions . . . . .	5
4.2. Device addressing . . . . .	6
4.3. General Message Types . . . . .	6
4.4. LoRaWAN MAC Frames . . . . .	6
5. SCHC over LoRaWAN . . . . .	6
5.1. Rule ID management . . . . .	6
5.2. IID computation . . . . .	6
5.3. Fragmentation . . . . .	6
5.3.1. Reliability options . . . . .	6
5.3.2. Supporting multiple window sizes . . . . .	11
5.3.3. Downlink fragment transmission . . . . .	11
5.3.4. SCHC behavior for devices in class A, B and C . . . . .	11
6. Security considerations . . . . .	11
7. Acknowledgements . . . . .	11
8. References . . . . .	11
8.1. Normative References . . . . .	11
8.2. Informative References . . . . .	12
Appendix A. Examples . . . . .	12
Appendix B. Note . . . . .	12
Authors' Addresses . . . . .	12

## 1. Introduction

The Static Context Header Compression (SCHC) specification [I-D.ietf-lpwan-ipv6-static-context-hc] describes generic header compression and fragmentation techniques that can be used on all LPWAN (Low Power Wide Area Networks) technologies defined in

[I-D.ietf-lpwan-overview]. Even though those technologies share a great number of common features like start-oriented topologies, network architecture, devices with mostly quite predictable communications, etc; they do have some slight differences in respect of payload sizes, reactivity, etc.

SCHC gives a generic framework that enables those devices to communicate with other Internet networks. However, for efficient performance, some parameters and modes of operation need to be set appropriately for each of the LPWAN technologies.

This document describes the efficient parameters and modes of operation when SCHC is used over LoRaWAN networks.

## 2. Terminology

This section defines the terminology and acronyms used in this document. For all other definitions, please look up the SCHC specification [I-D.ietf-lpwan-ipv6-static-context-hc].

- o DevEUI: an IEEE EUI-64 identifier used to identify the device during the procedure while joining the network (Join Procedure)
- o DevAddr: a 32-bit non-unique identifier assigned to a device statically or dynamically after a Join Procedure (depending on the activation mode)
- o TBD: all significant LoRaWAN-related terms.

## 3. Static Context Header Compression Overview

This section contains a short overview of Static Context Header Compression (SCHC). For a detailed description, refer to the full specification [I-D.ietf-lpwan-ipv6-static-context-hc].

Static Context Header Compression (SCHC) avoids context synchronization, which is the most bandwidth-consuming operation in other header compression mechanisms such as RoHC [RFC5795]. Based on the fact that the nature of data flows is highly predictable in LPWAN networks, some static contexts may be stored on the Device (Dev). The contexts must be stored in both ends, and it can either be learned by a provisioning protocol or by out of band means or it can be pre-provisioned, etc. The way the context is learned on both sides is out of the scope of this document.

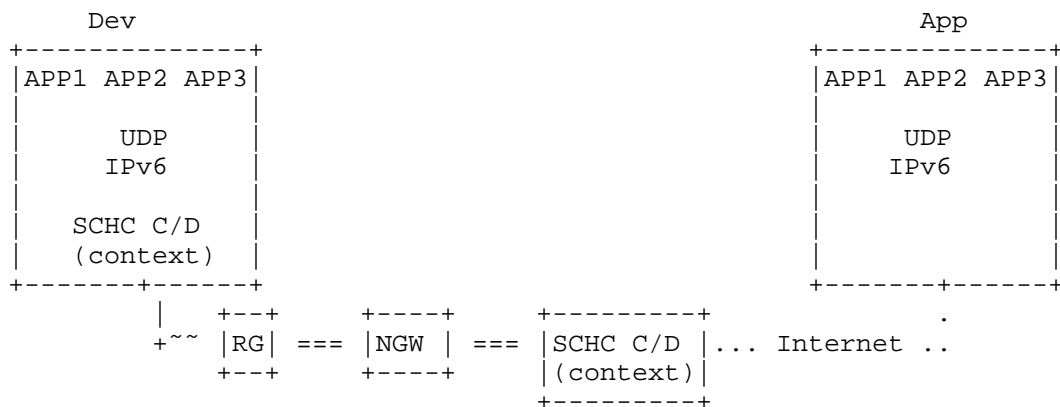


Figure 1: Architecture

Figure 1 represents the architecture for compression/decompression, it is based on [I-D.ietf-lpwan-overview] terminology. The Device is sending applications flows using IPv6 or IPv6/UDP protocols. These flows are compressed by an Static Context Header Compression Compressor/Decompressor (SCHC C/D) to reduce headers size. Resulting information is sent on a layer two (L2) frame to a LPWAN Radio Network (RG) which forwards the frame to a Network Gateway (NGW). The NGW sends the data to a SCHC C/D for decompression which shares the same rules with the Dev. The SCHC C/D can be located on the Network Gateway (NGW) or in another place as long as a tunnel is established between the NGW and the SCHC C/D. The SCHC C/D in both sides must share the same set of Rules. After decompression, the packet can be sent on the Internet to one or several LPWAN Application Servers (App).

The SCHC C/D process is bidirectional, so the same principles can be applied in the other direction.

In a LoRaWAN network, the RG is called a Gateway, the NGW is Network Server, and the SCHC C/D can be embedded in different places, for example in the Network Server and/or the Application Server.

Next steps for this section: detailed overview of the LoRaWAN architecture and its mapping to the SCHC architecture.

#### 4. LoRaWAN Architecture

An overview of LoRaWAN [lora-alliance-spec] protocol and architecture is described in [I-D.ietf-lpwan-overview]. Mapping between the LPWAN architecture entities as described in

[I-D.ietf-lpwan-ipv6-static-context-hc] and the ones in [lora-alliance-spec] is as follows:

- o Devices (Dev) are the end-devices or hosts (e.g. sensors, actuators, etc.). There can be a very high density of devices per radio gateway. This entity maps to the LoRaWAN End-device.
- o The Radio Gateway (RGW), which is the end point of the constrained link. This entity maps to the LoRaWAN Gateway.
- o The Network Gateway (NGW) is the interconnection node between the Radio Gateway and the Internet. This entity maps to the LoRaWAN Network Server.
- o LPWAN-AAA Server, which controls the user authentication and the applications. This entity maps to the LoRaWAN Join Server.
- o Application Server (App). The same terminology is used in LoRaWAN.




Figure 1: LPWAN/LoRaWAN Architecture

SCHC C/D (Compressor/Decompressor) and SCHC Fragmentation are performed on the LoRaWAN End-device and the Application Server. While the point-to-point link between the End-device and the Application Server constitutes single IP hop, the ultimate end-point of the IP communication may be an Internet node beyond the Application Server. In other words, the LoRaWAN Application Server acts as the first hop IP router for the End-device. Note that the Application Server and Network Server may be co-located, which effectively turns the Network/Application Server into the first hop IP router.

#### 4.1. Device classes (A, B, C) and interactions

TBD



#### 4.2. Device addressing

TBD

#### 4.3. General Message Types

TBD

#### 4.4. LoRaWAN MAC Frames

TBD

### 5. SCHC over LoRaWAN

#### 5.1. Rule ID management

Rule ID can be stored and transported in the FPort field of the LoRaWAN MAC frame or as the first bytes of the payload.

TBD

#### 5.2. IID computation

TBD

#### 5.3. Fragmentation

TBD

##### 5.3.1. Reliability options

###### 5.3.1.1. Uplinks: From device to gateway

In that case the device is the fragmentation transmitter, and the SCHC gateway the fragmentation receiver.

- o SCHC fragmentation reliability mode : "ACK\_ALWAYS"
- o Window size: 8, the FCN field is encoded on 3 bits
- o DTag : 1bit. this field is used to clearly separate two consecutive fragmentation sessions. A LoRaWAN device cannot interleave several fragmented SCHC datagrams.
- o MIC calculation algorithm: CRC32 using 0xEDB88320 (i.e. the reverse representation of the polynomial used e.g. in the Ethernet standard [RFC3385])

- o Retransmission Timer and inactivity Timer: LoRaWAN devices do not implement a "retransmission timer". At the end of a window the ACK corresponding to this window is transmitted by the network gateway in the RX1 or RX2 receive slot of the device. If this ACK is not received the device sends an all-0 (or an all-1) fragment with no payload to request an ACK retransmission. The periodicity between retransmission of the all-0/all-1 fragments is device/application specific and may be different for each device (not specified). The gateway implements an "inactivity timer". The default recommended duration of this timer is 12h. This value is mainly driven by application requirements and may be changed.

RuleID	DTag	W	FCN	Payload
3 bits	1 bit	1 bit	3 bits	

Figure 2: All fragment except the last one. Header size is 8 bits.

RuleID	DTag	W	FCN	MIC	Payload
3 bits	1 bit	1 bit	3 bits	32 bits	

Figure 3: All-1 fragment detailed format for the last fragment. Header size is 8 bits.

The format of an all-0 or all-1 acknowledge is:

RuleID	DTag	W	Encoded bitmap	Padding (0s)
3 bits	1 bit	1 bit	up to 8 bits	0 to 3 bits

Figure 4: ACK format for All-0 windows. Header size is 1 or 2 bytes.

RuleID	DTag	W	C	Encoded bitmap (if C = 0)	Padding (0s)
3 bits	1 bit	1 bit	1 bit	up to 8 bits	0 to 2 bits

Figure 5: ACK format for All-1 windows. Header size is 1 or 2 bytes.

5.3.1.2. Downlinks: From gateway to device

In that case the device is the fragmentation receiver, and the SCHC gateway the fragmentation transmitter. The following fields are common to all devices.

- o SCHC fragmentation reliability mode : ACK\_ALWAYS
- o Window size : 1 , The FCN field is encoded on 1 bits
- o DTag : 1bit. This field is used to clearly separate two consecutive fragmentation sessions. A LoRaWAN device cannot interleave several fragmented SCHC datagrams.
- o MIC calculation algorithm: CRC32 using 0xEDB88320 (i.e. the reverse representation of the polynomial used e.g. in the Ethernet standard [RFC3385])
- o MAX\_ACK\_REQUESTS : 8

RuleID	DTag	W	FCN	Payload	Padding
-----	-----	-----	-----	-----	-----
3 bits	1 bit	1 bit	1 bits	X bytes	2 bits

Figure 6: All fragments but the last one. Header size is 6 bits.

RuleID	DTag	W	FCN	MIC	Payload	Padding
-----	-----	-----	-----	-----	-----	-----
3 bits	1 bit	1 bit	1 bits	32 bits	X bytes	2 bits

Figure 7: All-1 Fragment Detailed Format for the Last Fragment. Header size is 6 bits.

The format of an all-0 or all-1 acknowledge is:

RuleID	DTag	W	Encoded bitmap	Padding (0s)
-----	-----	-----	-----	-----
3 bits	1 bit	1 bit	1 bit	2 bits

Figure 8: ACK format for All-0 windows. Header size is 8 bits.

RuleID	DTag	W	C = 1	Padding (0s)	
+ -----	+ -----	+ -----	+ -----	+ -----	+ -----
3 bits	1 bit	1 bit	1 bit	2 bits	

Figure 9: ACK format for All-1 windows, MIC is correct. Header size is 8 bits.

RuleID	DTag	W	b'111	0xFF (all 1's)	
+ -----	+ -----	+ -----	+ -----	+ -----	+ -----
3 bits	1 bit	1 bit	3 bits	8 bits	

Figure 10: Receiver ABORT packet (following an all-1 packet with incorrect MIC). Header size is 16 bits.

Class A and classB&C device do not manage retransmissions and timers in the same way.

#### 5.3.1.2.1. Class A devices

Class A devices can only receive in an RX slot following the transmission of an uplink. Therefore there cannot be a concept of "retransmission timer" for a gateway talking to classA devices for downlink fragmentation.

The device replies with an ACK fragment to every single fragment received from the gateway (because the window size is 1). Following the reception of a FCN=0 fragment (fragment that is not the last fragment of the packet or ACK-request), the device MUST transmit the ACK fragment until it receives the fragment of the next window. The device shall transmit up to MAX\_ACK\_REQUESTS ACK fragments before aborting. The device should transmit those ACK as soon as possible while taking into consideration eventual local radio regulation on duty-cycle, to progress the fragmentation session as quickly as possible. The ACK bitmap is 1 bit long and is always 1.

Following the reception of a FCN=1 fragment (the last fragment of a datagram) and if the MIC is correct, the device shall transmit the ACK with the "MIC is correct" indicator bit set. This message might be lost therefore the gateway may request a retransmission of this ACK in the next downlink. The device SHALL keep this ACK message in memory until it receives a downlink from the gateway different from an ACK-request indicating that the gateway has received the ACK message.

Following the reception of a FCN=1 fragment (the last fragment of a datagram) and if the MIC is NOT correct, the device shall transmit a receiver-ABORT fragment. The device SHALL keep this ABORT message in memory until it receives a downlink from the gateway different from an ACK-request indicating that the gateway has received the ABORT message. The fragmentation receiver (device) does not implement retransmission timer and inactivity timer.

The fragmentation sender (the gateway) implements an inactivity timer with default duration 12 hours. Once a fragmentation session is started, if the gateway has not received any ACK or receiver-ABORT message 12 hours after the last message from the device was received, the gateway may flush the fragmentation context. For devices with very low transmission rates (example 1 packet a day in normal operation) , that duration may be extended, but this is application specific.

#### 5.3.1.3. Class B or C devices

Class B&C devices can receive in scheduled RX slots or in RX slots following the transmission of an uplink. The device replies with an ACK fragment to every single fragment received from the gateway (because the window size is 1). Following the reception of a FCN=0 fragment (fragment that is not the last fragment of the packet or ACK-request), the device MUST always transmit the corresponding ACK fragment even if that fragment has already been received. The ACK bitmap is 1 bit long and is always 1. If the gateway receives this ACK, it proceeds to send the next window fragment. If the retransmission timer elapses and the gateway has not received the ACK of the current window it retransmits the last fragment. The gateway tries retransmitting up to MAX\_ACK\_REQUESTS times before aborting.

Following the reception of a FCN=1 fragment (the last fragment of a datagram) and if the MIC is correct, the device shall transmit the ACK with the "MIC is correct" indicator bit set. If the gateway receives this ACK, the current fragmentation session has succeeded and its context can be cleared.

If the retransmission timer elapses and the gateway has not received the all-1 ACK it retransmits the last fragment with the payload (not an ACK-request without payload). The gateway tries retransmitting up to MAX\_ACK\_REQUESTS times before aborting.

The device SHALL keep the all-1 ACK message in memory until it receives a downlink from the gateway different from the last (FCN=1) fragment indicating that the gateway has received the ACK message. Following the reception of a FCN=1 fragment (the last fragment of a datagram) and if the MIC is NOT correct, the device shall transmit a

receiver-ABORT fragment. The retransmission timer is used by the gateway (the sender), the optimal value is very much application specific but here are some recommended default values. For classB devices, this timer trigger is a function of the periodicity of the classB ping slots. The recommended value is equal to 3 times the classB ping slot periodicity. (modify 128sec) For classC devices which are nearly constantly receiving, the recommended value is 30 seconds. This means that the device shall try to transmit the ACK within 30 seconds of the reception of each fragment. The inactivity timer is implemented by the device to flush the context in-case it receives nothing from the gateway over an extended period of time. The recommended value is 12 hours for both classB&C devices.

#### 5.3.2. Supporting multiple window sizes

TBD

#### 5.3.3. Downlink fragment transmission

TBD

#### 5.3.4. SCHC behavior for devices in class A, B and C

TBD

### 6. Security considerations

TBD

### 7. Acknowledgements

TBD

### 8. References

#### 8.1. Normative References

- [RFC3385] Sheinwald, D., Satran, J., Thaler, P., and V. Cavanna, "Internet Protocol Small Computer System Interface (iSCSI) Cyclic Redundancy Check (CRC)/Checksum Considerations", RFC 3385, DOI 10.17487/RFC3385, September 2002, <<https://www.rfc-editor.org/info/rfc3385>>.
- [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>>.

- [RFC5795] Sandlund, K., Pelletier, G., and L-E. Jonsson, "The RObust Header Compression (ROHC) Framework", RFC 5795, DOI 10.17487/RFC5795, March 2010, <<https://www.rfc-editor.org/info/rfc5795>>.
- [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>>.

## 8.2. Informative References

- [I-D.ietf-lpwan-ipv6-static-context-hc]  
Minaburo, A., Toutain, L., Gomez, C., and D. Barthel,  
"LPWAN Static Context Header Compression (SCHC) and  
fragmentation for IPv6 and UDP", draft-ietf-lpwan-ipv6-  
static-context-hc-16 (work in progress), June 2018.
- [I-D.ietf-lpwan-overview]  
Farrell, S., "LPWAN Overview", draft-ietf-lpwan-  
overview-10 (work in progress), February 2018.
- [lora-alliance-spec]  
Alliance, L., "LoRaWAN Specification Version V1.0.2",  
<[http://portal.lora-  
alliance.org/DesktopModules/Inventures\\_Document/  
FileDownload.aspx?ContentID=1398](http://portal.lora-alliance.org/DesktopModules/Inventures_Document/FileDownload.aspx?ContentID=1398)>.

Appendix A. Examples

Appendix B. Note

Authors' Addresses

Nicolas Sornin (editor)  
Semtech  
14 Chemin des Clos  
Meylan  
France

Email: [nsornin@semtech.com](mailto:nsornin@semtech.com)

Michael Coracin  
Semtech  
14 Chemin des Clos  
Meylan  
France

Email: mcoracin@semtech.com

Ivaylo Petrov  
Acklio  
2bis rue de la Chataigneraie  
35510 Cesson-Sevigne Cedex  
France

Email: ivaylo@ackl.io

Alper Yegin  
Actility  
.  
Paris, Paris  
France

Email: alper.yegin@actility.com

Julien Catalano  
Kerlink  
1 rue Jacqueline Auriol  
35235 Thorigne-Fouillard  
France

Email: j.catalano@kerlink.fr

Vincent AUDEBERT  
EDF R&D  
7 bd Gaspard Monge  
91120 PALAISEAU  
FRANCE

Email: vincent.audebert@edf.fr



lpwan Working Group  
Internet-Draft  
Intended status: Informational  
Expires: January 3, 2019

JC. Zuniga  
SIGFOX  
C. Gomez  
Universitat Politecnica de Catalunya  
L. Toutain  
IMT-Atlantique  
July 02, 2018

SCHC over Sigfox LPWAN  
draft-zuniga-lpwan-schc-over-sigfox-03

Abstract

The Static Context Header Compression (SCHC) specification describes a header compression scheme and fragmentation functionality for Low Power Wide Area Network (LPWAN) technologies. SCHC offers a great level of flexibility that can be tailored for different LPWAN technologies.

The present document provides the optimal parameters and modes of operation when SCHC is implemented over a Sigfox LPWAN.

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
2. Terminology . . . . .	3
3. Static Context Header Compression . . . . .	3
4. SCHC over Sigfox . . . . .	4
4.1. SCHC Rules . . . . .	4
4.2. Packet processing . . . . .	4
5. Fragmentation . . . . .	4
5.1. Fragmentation headers . . . . .	5
5.2. Uplink fragment transmissions . . . . .	5
5.2.1. Uplink No-ACK mode . . . . .	5
5.2.2. Uplink ACK-Always mode . . . . .	6
5.2.3. Uplink ACK-on-Error mode . . . . .	6
5.3. Downlink fragment transmissions . . . . .	6
6. Padding . . . . .	7
7. Security considerations . . . . .	8
8. Acknowledgements . . . . .	8
9. Informative References . . . . .	8
Authors' Addresses . . . . .	9

## 1. Introduction

The Static Context Header Compression (SCHC) specification [I-D.ietf-lpwan-ipv6-static-context-hc] defines a header compression scheme and fragmentation functionality that can be used on top of all the LPWAN systems defined in [I-D.ietf-lpwan-overview]. These LPWAN systems have similar characteristics such as star-oriented topologies, network architecture, connected devices with built-in applications, etc.

SCHC offers a great level of flexibility to accommodate all these LPWAN systems. Even though there are a great number of similarities between LPWAN technologies, some differences exist with respect to the transmission characteristics, payload sizes, etc. Hence, there are optimal parameters and modes of operation that can be used when SCHC is used on top of a specific LPWAN.

This document describes the optimal parameters and modes of operation when SCHC is implemented over a Sigfox LPWAN.

2. Terminology

The reader is assumed to be familiar with the terms and mechanisms defined in [I-D.ietf-lpwan-overview] and in [I-D.ietf-lpwan-ipv6-static-context-hc].

3. Static Context Header Compression

Static Context Header Compression (SCHC) avoids context synchronization because data flows are highly predictable in LPWAN networks. Contexts must be stored and configured on both ends. This can be done either by using a provisioning protocol, by out of band means, or by pre-provisioning them e.g. at manufacturing time. The way the contexts are configured and stored on both ends is out of the scope of this document.

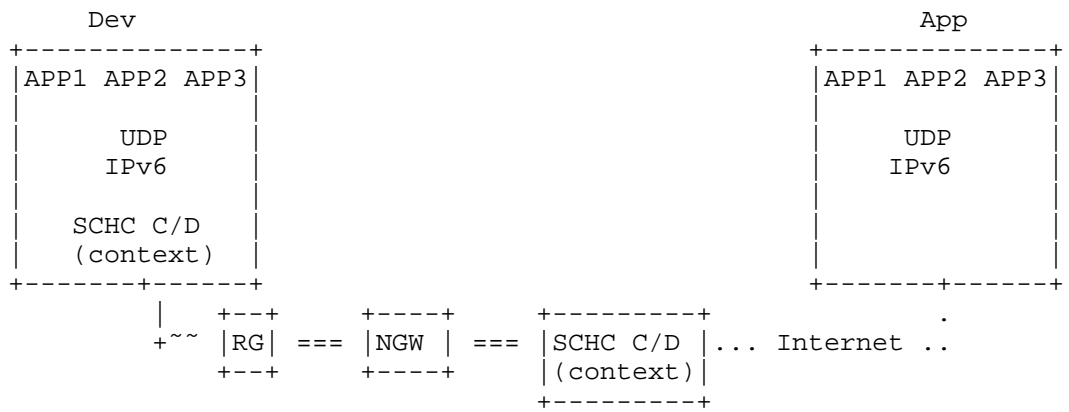


Figure 1: Architecture

Figure 1 represents the architecture for compression/decompression and fragmentation, which is based on [I-D.ietf-lpwan-overview] terminology.

The Device is sending applications flows that are compressed (and/or fragmented) by a Static Context Header Compression Compressor/Decompressor (SCHC C/D) to reduce headers size and/or fragment the packet. The resulting information is sent over a layer two (L2) frame to a LPWAN Radio Gateway (RG) which forwards the frame to a Network Gateway (NGW).

#### 4. SCHC over Sigfox

In the case of the global Sigfox network, RGs (or base stations) are distributed over the multiple countries where the Sigfox LPWAN service is provided. On the other hand, the NGW (or Cloud-based Core network) is a single entity that connects to all Sigfox base stations in the world.

Uplink transmissions occur in repetitions over different times and frequencies. Besides these time and frequency diversities, the Sigfox network also provides space diversity, as potentially an uplink message will be received by several base stations. Since all messages are self-contained and base stations forward them all back to the same Core network (NGW), multiple input copies can be combined at the NGW and hence provide for extra reliability based on the triple diversity.

The NGW communicates with the Network SCHC C/D for compression/decompression (and/or fragmentation/reassembly). The Network SCHC C/D shares the same set of rules as the Dev SCHC C/D. The Network SCHC C/D can be collocated with the NGW or in another place, as long as a tunnel is established between the NGW and the SCHC C/D. After decompression (and/or reassembly), the packet can be forwarded over the Internet to one (or several) LPWAN Application Server(s) (App).

The SCHC C/D process is bidirectional, so the same principles can be applied on both uplink and downlink.

##### 4.1. SCHC Rules

The RuleID MUST be sent at the beginning of the SCHC header. The total number of rules to be used affects directly the Rule ID field size, and therefore the total size of the fragmentation header. For this reason, it is recommended to keep the number of rules that are defined for a specific device to the minimum possible.

##### 4.2. Packet processing

TBD

#### 5. Fragmentation

The SCHC specification [I-D.ietf-lpwan-ipv6-static-context-hc] defines a generic fragmentation functionality that allows sending data packets larger than the maximum size of a Sigfox data frame. The functionality also defines a mechanism to send reliably multiple frames, by allowing to resend selectively any lost frames.

The SCHC fragmentation supports several modes of operation. These modes have different advantages and disadvantages depending on the specifics of the underlying LPWAN technology and Use Case. This section describes how the SCHC fragmentation functionality should optimally be implemented when used over a Sigfox LPWAN for the most typical use case applications.

### 5.1. Fragmentation headers

A list of fragmentation header fields, their sizes as well as recommended modes for SCHC fragmentation over Sigfox are provided in this section.

### 5.2. Uplink fragment transmissions

Uplink transmissions are completely asynchronous and can take place in any random frequency of the allowed uplink bandwidth allocation. Hence, devices can go to deep sleep mode, and then wake up and transmit whenever there is a need to send any information to the network. In that way, there is no need to perform any network attachment, synchronization, or other procedure before transmitting a data packet. All data packets are self contained with all the required information for the network to process them accordingly.

Since uplink transmissions occur asynchronously, an SCHC fragment can be transmitted at any given time by the Dev.

#### 5.2.1. Uplink No-ACK mode

No-ACK is RECOMMENDED to be used for transmitting short, non-critical packets that require fragmentation.

Fragmentation Header size: 8 bits

The recommended Rule ID size is: 3 bits

The recommended DTag size (T) is: 0 bits, as the number of available Rule IDs are sufficient to interleave fragmented packets.

Fragment Compressed Number (FCN) size (N): 4 bits

As per [REF SCHC], in the No-ACK mode the W (window) 1-bit field is not present.

When fragmentation is used to transport IP frames, the Message Integrity Check (MIC) size, M: TBD bits

The algorithm for computing the MIC field MUST be TBD.

### 5.2.2. Uplink ACK-Always mode

TBD

### 5.2.3. Uplink ACK-on-Error mode

ACK-on-Error is RECOMMENDED for larger packets, since it leads to a reduced number of ACKs to be sent in the lower capacity downlink channel.

The recommended Fragmentation Header size is: 8 bits

The recommended Rule ID size is: 3 bits.

The recommended DTag size (T) is: 0 bits, as the number of available Rule IDs are sufficient to interleave fragmented packets.

Fragment Compressed Number (FCN) size (N): 4 bits.

As per [REF SCHC], in the ACK-on-Error mode the Window (W) 1-bit field must be present.

For the ACK-on-Error fragmentation mode(s), a single window size is RECOMMENDED.

The value of MAX\_ACK\_REQUESTS SHOULD be 2, and the value of MAX\_WIND\_FCN SHOULD be 14 (which allows a maximum window size with 15 fragments).

When fragmentation is used to transport IP frames, the Message Integrity Check (MIC) size, M: TBD bits

The algorithm for computing the MIC field MUST be TBD.

## 5.3. Downlink fragment transmissions

In some LPWAN technologies, as part of energy-saving techniques, downlink transmission is only possible immediately after an uplink transmission. This allows the device to go in a very deep sleep mode and preserve battery, without the need to listen to any information from the network. This is the case for Sigfox-enabled devices, which can only listen to downlink communications after performing an uplink transmission.

When there are multiple fragments to be transmitted in the downlink, an uplink message is required to trigger the downlink communication. In order to avoid potentially high delay for fragmented datagram transmission in the downlink, the fragment receiver MAY perform an

uplink transmission as soon as possible after reception of a fragment that is not the last one. Such uplink transmission MAY be triggered by sending a SCHC message, such as an ACK. In this sense, ACK-Always is the preferred fragmentation mode for downlink communications.

For downlink fragment transmission, the ACK-Always mode MUST be supported.

The recommended Fragmentation Header size is: 8 bits

The recommended Rule ID size is: 3 bits.

The recommended DTag size (T) is: 0 bits, as the number of available Rule IDs are sufficient to interleave fragmented packets.

Fragment Compressed Number (FCN) size (N): 4 bits.

As per [REF SCHC], in the ACK-on-Error mode the Window (W) 1-bit field must be present.

For the ACK-Always fragmentation mode(s), a single window size is RECOMMENDED.

The value of MAX\_ACK\_REQUESTS SHOULD be 2, and the value of MAX\_WIND\_FCN SHOULD be 14 (which allows a maximum window size with 15 fragments).

When fragmentation is used to transport IP frames, the Message Integrity Check (MIC) size, M: TBD bits

The algorithm for computing the MIC field MUST be TBD.

Sigfox downlink frames have a fixed length of 8 bytes, which means that default SCHC algorithm for padding cannot be used. Therefore, the 3 last bits of the fragmentation header are used to indicate in bytes the size of the padding. A size of 000 means that the full remaining frame is used to carry payload, a value of 001 indicates that the last byte contains padding, and so on.

## 6. Padding

The Sigfox payload fields have different characteristics in uplink and downlink.

Uplink frames can contain a payload from 0 to 96 bits (i.e. 12 bytes). The radio protocol allows sending zero bits or one single bit of information for binary applications (e.g. status). However, for 2 or more bits of payload it is required to add padding to the

next integer number of bytes. The reason for this flexibility is to optimize transmission time and hence save battery consumption at the device.

Downlink frames on the other hand have a fixed length. The payload length must be 64 bits (i.e. 8 bytes). Hence, if less information bits are to be transmitted padding would be necessary and it should be performed as described in the previous section.

## 7. Security considerations

The radio protocol authenticates and ensures the integrity of each message. This is achieved by using a unique device ID and an AES-128 based message authentication code, ensuring that the message has been generated and sent by the device with the ID claimed in the message.

Application data can be encrypted at the application level or not, depending on the criticality of the use case, to provide a balance between cost and effort vs. risk. AES-128 in counter mode is used for encryption. Cryptographic keys are independent for each device. These keys are associated with the device ID and separate integrity and confidentiality keys are pre-provisioned. A confidentiality key is only provisioned if confidentiality is to be used.

The radio protocol has protections against reply attacks, and the cloud-based core network provides firewalling protection against undesired incoming communications.

## 8. Acknowledgements

Carles Gomez has been funded in part by the ERDF and the Spanish Government through project TEC2016-79988-P.

## 9. Informative References

[I-D.ietf-lpwan-ipv6-static-context-hc]  
Minaburo, A., Toutain, L., and C. Gomez, "LPWAN Static Context Header Compression (SCHC) and fragmentation for IPv6 and UDP", draft-ietf-lpwan-ipv6-static-context-hc-07 (work in progress), October 2017.

[I-D.ietf-lpwan-overview]  
Farrell, S., "LPWAN Overview", draft-ietf-lpwan-overview-07 (work in progress), October 2017.



Authors' Addresses

Juan Carlos Zuniga  
SIGFOX  
425 rue Jean Rostand  
Labege 31670  
France

Email: [JuanCarlos.Zuniga@sigfox.com](mailto:JuanCarlos.Zuniga@sigfox.com)  
URI: <http://www.sigfox.com/>

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

Email: [carlesgo@entel.upc.edu](mailto:carlesgo@entel.upc.edu)

Laurent Toutain  
IMT-Atlantique  
2 rue de la Chataigneraie  
CS 17607  
35576 Cesson-Sevigne Cedex  
France

Email: [Laurent.Toutain@imt-atlantique.fr](mailto:Laurent.Toutain@imt-atlantique.fr)