

lpwan
Internet-Draft
Intended status: Informational
Expires: May 1, 2017

S. Farrell
Trinity College Dublin
A. Yegin
Actility
October 28, 2016

LoRaWAN Overview
draft-farrell-lpwan-lora-overview-01

Abstract

Low Power Wide Area Networks (LPWAN) are wireless technologies covering different Internet of Things (IoT) applications. The common characteristics for LPWANs are large coverage, low bandwidth, small packet and application layer data sizes and long battery life operation. One of these technologies is LoRaWAN developed by the LoRa Alliance. LoRaWAN targets key requirements of the Internet of things such as secure bi-directional communication, mobility and localization services. This memo is an informational overview of LoRaWAN and gives the principal characteristics of this technology in order to help with the IETF work for providing IPv6 connectivity over LoRaWAN along with other 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 May 1, 2017.

Copyright Notice

Copyright (c) 2016 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. Radio Spectrum	4
4. MAC Layer	6
5. Names and Addressing	8
6. Security Considerations	10
6.1. Payload Encryption and Data Integrity	10
6.2. Key Derivation	10
7. IANA Considerations	11
8. Acknowledgements	11
9. Contributors	11
10. Informative References	12
Authors' Addresses	12

1. Introduction

LoRaWAN is a wireless technology for long-range low-power low-data-rate applications developed by the LoRa Alliance, a membership consortium. <https://www.lora-alliance.org/> 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. All communication is generally bi-directional, although uplink communication from end-devices to the network server are favoured in terms of overall bandwidth availability.

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.

To maximize both battery life of end-devices and overall network capacity, 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, using any of the currently available data rates.

This memo provides an overview of the LoRaWAN technology for the Internet community, but the definitive specification [LoRaSpec] is that produced by the LoRa Alliance. This draft is based on version 1.0.2 of the LoRa specification. (Note that version 1.0.2 is expected to be published in a few weeks. We will update this draft when that has happened. For now, version 1.0 is available at [LoRaSpec1.0])

2. Terminology

This section introduces some LoRaWAN terms. 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 Uplink message: refers to communications from end-device to network server or application via one or more gateways.
- o Downlink message: refers to communications from 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.
- o Classes A, B and C define different device capabilities and modes of operation for end-devices. End-devices can transmit uplink messages at any time in any mode of operation (so long as e.g., ISM band restrictions are honoured). An end-device in Class A can only receive downlink messages at predetermined timeslots after each uplink message transmission. Class B allows the end-device to receive downlink messages at periodically scheduled timeslots. Class C allows receipt of downlink messages at anytime. Class selection is based on the end-devices' application use case and its power supply. (While Classes B and C are not further described here, readers may have seen those terms elsewhere so we include them for clarity.)

3. Radio Spectrum

LoRaWAN radios make use of 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.

As with other LPWAN radio technologies, LoRaWAN end-devices respect the frequency, power and maximum transmit duty cycle requirements for the sub-band imposed by local regulators. In most cases, this means an end-device is only transmitting for 1% of the time, as specified by ISM band regulations. And in some cases the LoRaWAN specification calls for end-devices to transmit less often than is called for by the ISM band regulations in order to avoid congestion.

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. The frequencies and data rate chosen for the first of these receive windows match those used for the transmit window. The frequency and data-rate for the second receive window are configurable. If a downlink message preamble is detected during a receive window, then the end-device keeps the radio on in order to receive the frame.

End-devices can only transmit a subsequent uplink frame after the end of the associated receive windows. When a device joins a LoRaWAN

network (see Section 4 for details), 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: 0.3-5 kbps

Table 1: Default settings for EU868MHz 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 stringent 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 that can send the 23 byte Join-request message for the join procedure.

4. MAC Layer

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 join procedure.

The header also distinguishes the uplink/downlink directions and whether or not an acknowledgement ("confirmation") is required from the peer.

All payloads are encrypted and ciphertxts are protected with a cryptographic Message Integrity Check (MIC) - see Section 6 for details.

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").

LoRaWAN end-devices can choose various different data rates from a menu of available rates (dependent on the frequencies in use). It is however, recommended that end-devices set the Adaptive Data Rate ("ADR") bit in the MAC layer which is a signal that the network should control the data rate (via MAC commands to the end-device). The network can also assert the ADR bit and control data rates at its discretion. The goal is to ensure minimal on-time for radios whilst increasing throughput and reliability when possible. Other things being equal, the effect should be that end-devices closer to a gateway can successfully use higher data rates, whereas end-devices further from all gateways still receive connectivity though at a lower data rate.

Data rate changes can be validated via a scheme of acks from the network with a fall-back to lower rates in the event that downlink acks go missing.

There are 16 (or 32) bit frame counters maintained in each direction that are incremented on each transmission (but not re-transmissions) that are not re-used for a given key. When the device supports a 32 bit counter, then only the least significant 16 bits are sent in the MAC header, but all 32 bits are used in cryptographic operations. (If an end-device only supports a 16 bit counter internally, then the topmost 16 bits are set to zero.)

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).

The network server can also inform an end-device about channel assignments (mid-point frequencies and data rates). Of course, these must also remain within the bands assigned by local regulation.

5. Names and Addressing

A LoRaWAN network has a short network identifier ("NwkID") which is a seven bit value. A private network (common for LoRaWAN) can use the value zero. If a network wishes to support "foreign" end-devices then the NwkID needs to be registered with the LoRA Alliance, in which case the NwkID is the seven least significant bits of a registered 24-bit NetID. (Note however, that the methods for "roaming" are currently being enhanced within the LoRA Alliance, so the situation here is somewhat fluid.)

In order to operate nominally on a LoRaWAN network, a device needs a 32-bit device address, which is the concatenation of the NwkID and a 25-bit device-specific network address that is assigned when the device "joins" the network (see below for the join procedure) or that is pre-provisioned into the device.

End-devices are assumed to work with one or a quite limited number of applications, which matches most LoRaWAN use-cases. The applications are 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 artefacts (port=0), the NwkSKey, and another for protecting application layer traffic, the AppSKey. Both keys are used for 128 bit AES cryptographic operations. (See Section 6 for details.)

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 summarises these values.

Value	Description
DevAddr	DevAddr (32-bits) = NwkId (7-bits) + device-specific network address (25 bits)
AppEUI	IEEE EUI64 naming the application
NwkSKey	128 bit network session key for use with AES
AppSKey	128 bit application session key for use with AES

Table 3: Values required for nominal operation

As an alternative, end-devices can use the LoRaWAN join procedure 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. In addition to the AppEUI, end-devices using the join procedure need to also know 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 summarises 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.

There is some special handling related to which channels to use and for multiple transmissions for the join-request which is intended to ensure a successful join in as many cases as possible. Join-request and Join-accept messages also include some random values (nonces) to both provide some replay protection and to help ensure the session keys are unique per run of the join procedure. If a Join-request fails validation, then no Join-accept message (indeed no message at all) is returned to the end-device. For example, if an end-device is factory-reset then it should end up in a state in which it can re-do the join procedure.

6. Security Considerations

In this section we describe the use of cryptography in LoRaWAN. This section is not intended as a full specification but to be sufficient so that future IETF specifications can encompass the required security considerations. The emphasis is on describing the externally visible characteristics of LoRaWAN.

6.1. Payload Encryption and Data Integrity

All payloads are encrypted and have data integrity. Frame options (used for 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. Since MAC commands may be sent as options and not only as payload, any values sent in that manner are visible to a passive attacker but are not malleable for an active attacker due to the use of the MIC.

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 Message Integrity Code (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.

6.2. Key Derivation

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.

7. IANA Considerations

There are no IANA considerations related to this memo.

8. Acknowledgements

The authors re-used some text from [I-D.vilajosana-lpwan-lora-hc]

Stephen Farrell's work on this memo was supported by the Science Foundation Ireleand funded CONNECT centre <<https://connectcentre.ie/>>.

9. Contributors

The following members of the LoRa Alliance reviewed this draft and contributed (much more than SF) to the definition of LoRaWAN.

Name, Affiliation, email (optional)

Chun-Yeow Yeoh, VADS LYFE SDN BHD, yeow@tmrnd.com.my

Olivier Hersent, Actility, olivier.hersent@actility.com

Dave Kjendal, Senet Inc, dkjendal@senetco.com

Paul Duffy, Cisco, paduffy@cisco.com

Joachim Ernst, Swisscom Broadcast Ltd, joachim.ernst@swisscom.com

Nicolas Sornin, Semtech, nsornin@semtech.com

Phillippe Christin, Orange, philippe.christin@orange.com

10. Informative References

[I-D.vilajosana-lpwan-lora-hc]

Vilajosana, X., Dohler, M., and A. Yegin, "Transmission of IPv6 Packets over LoRaWAN", draft-vilajosana-lpwan-lora-hc-00 (work in progress), July 2016.

[LoRaSpec]

LoRa Alliance, "LoRaWAN Specification Version V1.0.2", Nov 2016, <URL TBD>.

[LoRaSpec1.0]

LoRa Alliance, "LoRaWAN Specification Version V1.0", Jan 2015, <<https://www.lora-alliance.org/portals/0/specs/LoRaWAN%20Specification%201R0.pdf>>.

Authors' Addresses

Stephen Farrell
Trinity College Dublin
Dublin 2
Ireland

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

Alper Yegin
Actility
Paris, Paris
FR

Email: alper.yegin@actility.com

lpwan
Internet-Draft
Intended status: Informational
Expires: May 4, 2017

S. Farrell, Ed.
Trinity College Dublin
October 31, 2016

LPWAN Overview
draft-farrell-lpwan-overview-04

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 May 4, 2017.

Copyright Notice

Copyright (c) 2016 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. 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) 10
 - 2.2.1. Provenance and Documents 10
 - 2.2.2. Characteristics 11
 - 2.3. SIGFOX 14
 - 2.3.1. Provenance and Documents 15
 - 2.3.2. Characteristics 15
 - 2.4. Wi-SUN Alliance Field Area Network (FAN) 19
 - 2.4.1. Provenance and Documents 19
 - 2.4.2. Characteristics 20
- 3. Generic Terminology 22
- 4. Gap Analysis 23
 - 4.1. Naive application of IPv6 23
 - 4.2. 6LOWPAN 24
 - 4.2.1. Header Compression 24
 - 4.2.2. Address Autoconfiguration 25
 - 4.2.3. Fragmentation 25
 - 4.2.4. Neighbor Discovery 25
 - 4.3. 6lo 26
 - 4.4. 6tisch 26
 - 4.5. RoHC 27
 - 4.6. ROLL 27
 - 4.7. CoAP 27
 - 4.8. Mobility 28
 - 4.9. DNS and LPWAN 28
- 5. Security Considerations 28
- 6. IANA Considerations 29
- 7. Contributors 29
- 8. Acknowledgements 32
- 9. Informative References 32
- Author's Address 35

1. Introduction

[[Ed: Editor comments/queries are in double square brackets like this. Note that the eventual fate of this draft is a topic for the WG to consider - it might end up as a useful RFC, or it might be best maintained as a draft only until its utility has dissipated. FWIW, the editor doesn't mind what outcome the WG choose.]]

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 low-cost, low-throughput devices at very low-cost and with very-low power consumption, so that even battery-powered devices can be deployed for years. 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.

Existing pilot deployments have shown huge potential and created much industrial interest in these technologies. As of today, [[Ed: with the possible exception of Wi-SUN devices?]] essentially no LPWAN devices 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 LPWAN WG is to 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. Discussion of this document should take place on the lp-wan@ietf.org list.

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. [[Ed: I assume that's the right target here. Please comment if you disagree.]]

[[Ed: the goal here is 2-3 pages per technology. If there's much more needed then we could add appendices I guess depending on what text the WG find useful to include.]]

[[Ed: A lot of the radio frequency related details below could disappear I think - for the purposes of this WG, I think a lot of that is extraneous detail. Haven't yet done that though, in case I'm

missing something. It might also further imbalance the level of description of the different technologies, to the extent that the WG care explicitly about that.]]

2.1. LoRaWAN

[[Ed: Text here is from [I-D.farrell-lpwan-lora-overview]]]

2.1.1. Provenance and Documents

LoRaWAN is a 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. (Version 1.0.2 is expected to be published in a few weeks. We will wmen that has happened. For now, version 1.0 is available at [LoRaSpec1.0])

2.1.2. Characteristics

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, using any of the currently available data rates.

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. All communication is generally bi-directional, although uplink communication from end-devices to the network server are favoured in terms of overall bandwidth availability.

Figure 1 shows the entities involved in a LoRaWAN network.

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: 0.3-5 kbps

Table 1: Default settings for EU868MHz 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 stringent 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 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 join procedure.

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

it's 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).

A LoRaWAN network has a short network identifier ("NwkID") which is a seven bit value. A private network (common for LoRaWAN) can use the value zero. If a network wishes to support "foreign" end-devices then the NwkID needs to be registered with the LoRA Alliance, in which case the NwkID is the seven least significant bits of a registered 24-bit NetID. (Note however, that the methods for "roaming" are currently being enhanced within the LoRA Alliance, so the situation here is somewhat fluid.)

In order to operate nominally on a LoRaWAN network, a device needs a 32-bit device address, which is the concatenation of the NwkID and a 25-bit device-specific network address that is assigned when the device "joins" the network (see below for the join procedure) or that is pre-provisioned into the device.

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 artefacts (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 summarises these values.

Value	Description
DevAddr	DevAddr (32-bits) = NwkId (7-bits) + device-specific network address (25 bits)
AppEUI	IEEE EUI64 naming the application
NwkSKey	128 bit network session key for use with AES
AppSKey	128 bit application session key for use with AES

Table 3: Values required for nominal operation

As an alternative, end-devices can use the LoRaWAN join procedure 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 summarises 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 MIC.

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 Message Integrity Code (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)

[[Ed: Text here is from [I-D.ratilainen-lpwan-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, but further enhancements for NB-IoT are worked on in the following releases, for example in the form of multicast support. 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, RLC, PDCP and RRC protocol layers respectively.

2.2.2. Characteristics

[[Ed: Not clear what minimum/worst-case MTU might be. There are many 3GPP acronyms/terms to eliminate or explain.]]

Specific targets for NB-IoT include: Less than 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 size MTU of 1600 bytes. As the name suggests, NB-IoT uses narrowbands with the 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. On uplink multi-tone SC-FDMA is used with 15 kHz tone spacing or as a special case of SC-FDMA single tone with either 15kHz or 3.75 kHz tone spacing may be used.

NB-IoT can be deployed in three ways. In-band deployment means that the narrowband is multiplexed within normal LTE carrier. In Guard-band deployment the narrowband uses the unused resource blocks between two adjacent LTE carriers. Also standalone deployment is supported, where the narrowband can be located alone in dedicated spectrum, which makes it possible for example to refarm the GSM carrier at 850/900 MHz for NB-IoT. All three deployment modes are meant to be used in licensed 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.

For signaling optimization, two options are introduced in addition to legacy 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 from Mobility Management Entity (MME) in core network to the UE without interaction from base station. This means there are no Access Stratum security or header compression, as the Access Stratum is bypassed, and only limited RRC procedures.

The RRC Suspend/Resume procedures reduce the signaling overhead required for UE state transition from Idle to Connected mode in order to have a user plane transaction with the network and back to Idle state by reducing the signaling messages required compared to legacy operation

With extended DRX the RRC Connected mode DRX cycle is up to 10.24 seconds and in RRC Idle the DRX cycle can be up to 3 hours.

NB-IoT has no channel access restrictions allowing up to a 100% duty-cycle.

3GPP access security is specified in [TGPP33203].




Figure 3: 3GPP network architecture

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 other 3GPP technologies.

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

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 illustration of 3GPP radio protocol architecture can be seen from Figure 4.




Figure 4: 3GPP radio protocol architecture

The radio protocol architecture of NB-IoT (and LTE) is separated into control plane and user plane. 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 EPC, passing transparently through radio network. It is responsible for authentication, security control, mobility management and bearer management.

Access Stratum (AS) is the functional layer below NAS, and in 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. 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 singlecast data. In RRC_Connected the UE has a connection to the eNodeB, the network knows the UE location on cell level and the UE may receive and transmit singlecast data. RRC_Connected mode is established, when the UE is expected to be active in the network, to transmit or receive data. Connection is released, switching to RRC_Idle, when there is no traffic to save the UE battery 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, while the UE is in RRC_Idle transparently to the eNodeB.

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 consists of 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. Lower layers in user plane are similarly RLC, MAC and physical layer performing tasks mentioned above.

Under worst-case conditions, NB-IoT may achieve data rate of roughly 200 bps. For downlink with 164 dB coupling loss, NB-IoT may achieve higher data rates, 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 higher maximum coupling loss than 170 dB with very low bit rates.

2.3. SIGFOX

[[Ed: Text here is from
[I-D.zuniga-lpwan-sigfox-system-description].]]

2.3.1. Provenance and Documents

The SIGFOX LPWAN is in line with the terminology and specifications being defined by the ETSI ERM TG28 Low Throughput Networks (LTN) group [etsi_ltn]. As of today, SIGFOX's network has been fully deployed in 6 countries, with ongoing deployments on 18 other countries, in total a geography containing 397M 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. The capacity of a SIGFOX base station mainly depends on the number of messages generated by the devices, and not on the number of devices. The battery life of devices also depends on the number of messages generated by the device, but it is important to keep in mind that these devices are designed to last several years, some of them even buried underground. The coverage of the cell also depends on the link budget and on the type of deployment (urban, rural, etc.), which can vary from sending less than one message per device per day to about ten messages per device per day.

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 LTN 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 communications are possible with the UNB solution. 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 in the USA)
- o Uplink baud rate: 100 baud (600 baud in the USA)
- 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

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 maximum mean transmission time 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 in the USA)
- o Link budget: 153 dB (or better)
- o Central frequency accuracy: Centre frequency of downlink transmission are set by the network according to the corresponding uplink transmission.

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 querying the network for existing data from the device.

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 LTN network transmits the downlink message for a given device during the reception window. The LTN network selects the BS for transmitting the corresponding downlink message.

Uplink and downlink transmissions are unbalanced due to the regulatory constraints on the ISM bands. Under the strictest regulations, the system can allow a maximum of 140 uplink messages and 4 downlink messages per device. These restrictions can be slightly relaxed depending on system conditions and the specific regulatory domain of operation.




Figure 7: ETSI LTN architecture

Figure 7 depicts the different elements of the SIGFOX architecture.

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

The architecture consists of a single 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 or End Points (EPs) are the objects that communicate application data between local device applications (DAs) and network applications (NAs).

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

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

The radio protocol provides mechanisms to authenticate and ensure integrity of the message. This is achieved by using a unique device ID and a message authentication code, which allow ensuring that the message has been generated and sent by the device with the ID claimed in the message.

Security keys are independent for each device. These keys are associated with the device ID and they are pre-provisioned. Application data can be encrypted by the application provider.

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

[[Ed: Text here is via personal communication from Bob Heile (bheile@ieee.org) and was authored by Bob and Sum Chin Sean. Many references to specifications are still needed here.]]

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 [[Ed: reference needed!]] is an IPv6 frequency hopping 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. [[Ed: details welcome.]]

2.4.2. Characteristics

[[Ed: this really needs the references.]] The FAN profile is based on various open standards in IETF, IEEE802 and ANSI/TIA for low power and lossy networks. The FAN profile specification provides an application-independent IPv6-based transport service for both connectionless (i.e. UDP) and connection-oriented (i.e. TCP) services. There are two possible methods for establishing the IPv6 packet routing: mandatory Routing Protocol for Low-Power and Lossy Networks (RPL) at the Network layer or optional Multi-Hop Delivery Service (MHDS) 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 defined in RFC2460 with 6LoWPAN adaptation as defined in RC4944 and RFC6282. Additionally, ICMPv6 as defined in RFC4443 is used for control plane in 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, such as IEEE802.15.4k.

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. 802.11i Group Key Management 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

[[Ed: Text here is from [I-D.minaburo-lpwan-gap-analysis].]]

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 The Host, which 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 hosts 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 AAA Server, which controls the user authentication, the applications. It is known as: Join-Server, Home Subscriber Server or Registration Authority. [[Ed: I'm not clear that AAA server is the right generic term here.]]
- o At last we have the Application Server, known also as Packet Data Node Gateway or Network Application.

Function/ Technology	LORAWAN	NB-IOT	SIGFOX	IETF
Sensor, Actuator, device, object	End Device	User Equipment	End Point	Thing (HOST)
Transceiver Antenna	Gateway	Evolved Node B	Base Station	RADIO GATEWAY
Server	Network Server	Serving- Gateway	Service Center	Network Gateway (ROUTER)
Security Server	Join Server	Home Subscriber Server	Registration Authority	AAA SERVER
Application	Application Server	Packet Data Node Gateway	Network Application	APPLICATION SERVER

Figure 8: LPWAN Architecture Terminology




Figure 9: LPWAN Architecture

4. Gap Analysis

[[Ed: Text here is from [I-D.minaburo-lpwan-gap-analysis].]]

4.1. Naive application of IPv6

IPv6 [RFC2460] 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 would 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 needs 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.

[[Ed: other things to maybe mention: IPsec, DHCPv6, anything with even 1 regular RTT needed, e.g. DNS.]]

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 characterised by device constraints (in terms of processing capacity, memory, and energy availability), and specially, link constraints, such as:

- o very low 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 optimisations 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.

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 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 or IEEE 802.11ah. These technologies are similar in several aspects to IEEE 802.15.4, which was the original 6LoWPAN target technology. [[Ed: refs?]]

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. [[Ed: I'm not sure that conclusion follows from the information provided in this section - is more needed?.]]

4.4. 6tisch

The 6tisch solution is dedicated to mesh networks that operate using 802.15.4e MAC with a deterministic slotted channel. The TSCH [[Ed: expand on 1st use]] 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

RoHC [[Ed: expand on 1st use]] header compression mechanisms were defined for point to point multimedia channels, to reduce the header overhead of RTP flows. RoHC can also reduce the overhead of IPv4 or IPv6 or UDP headers. It is based on shared context which does not require any state but compressed packets are not routable. The context is initialised at the beginning of the communication or when it [[Ed: which "it"?]] is lost. The compression is managed using a sequence number (SN) which is encoded using a windowing algorithm allowing for reduction of the SN to 4 bits instead of 2 bytes. [[Ed: is that the 2 bytes as per 6lowPAN?]] But this window needs to be updated each 15 packets which implies larger headers. When RoHC is used we talk about an average header compression size to give the performance of compression. For example, RoHC starts sending bigger packets than the original (52 bytes) to reduce the header up to 4 bytes (it stays here only for 15 packets, which correspond to the window size). Each time the context is lost or needs to be synchronised, packets of about 15 to 43 bytes are sent. [[Ed: the above isn't that cleaar to me.]]

RoHC is not adapted to the constrained nodes of the LPWAN networks: it does not take into account the energy limitations and the transmission rate, and context is synchronised during the transmission, which does not allow a better compression. [[Ed: this seems to conflict a bit with what was said about 6tisch which puzzled me.]]

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 acknowledgements may allow the reduction of L2 acknowledgements. 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 proves 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

LPWANs 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. The mobility will imply most of the time a group of devices, which represent a network itself. The mobility concerns more the gateway than the devices.

NEMO [[Ed: refs?]] Mobility solutions may be used in the case where some hosts belonging to the same Network gateway will move from one point to another and that they are not aware of this mobility.

4.9. DNS and LPWAN

The purpose of the DNS is to enable applications to name things that have a global unique name. Lots of protocols are using DNS to identify the objects, especially REST and applications using CoAP. Therefore, hosts (things), or the named services they use, should be registered in DNS. DNS is probably a good topic of research for LPWAN technologies, while the matching of the name and the IP information can be used to configure the LPWAN devices. [[Ed: I'm not sure what that last bit means.]]

5. Security Considerations

[[Ed: be good to add stuff here about a) privacy and b) difficulties with getting current security protocols to work in this context. For a) maybe try find nice illustrations, e.g. extremecom instrumeted-igloo traces (temperature change allowing one to infer when someone took a pee:-). For b) things like IPsec/(D)TLS/OCSP and NTP to work in these environments. Not sure how much of that is known or useful for the WG. Probably worth noting the IAB statement on confidentiality and to ponder the impact of more than one layer of encryption in this context. Text below is basically from the "gaps" draft.]]

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.

6. IANA Considerations

There are no IANA considerations related to this memo.

7. Contributors

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

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

Email: 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

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

Email: ana@ackl.io

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

Email: josep.paradells@entel.upc.edu

Benoit Ponsard
SIGFOX
425 rue Jean Rostand
Labege 31670
France

Email: Benoit.Ponsard@sigfox.com
URI: <http://www.sigfox.com/>

Antti Ratilainen
Ericsson
Hirsalantie 11
Jorvas 02420
Finland

Email: 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

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

Alper Yegin
Actility
Paris, Paris
FR

Email: 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. Acknowledgements

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

In addition to the contributors above, thanks are due to Jiazi Yi, [your name here] for comments.

Stephen Farrell's work on this memo was supported by the Science Foundation Ireland funded CONNECT centre <<https://connectcentre.ie/>>.

9. Informative References

- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <<http://www.rfc-editor.org/info/rfc2460>>.
- [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, <<http://www.rfc-editor.org/info/rfc2904>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<http://www.rfc-editor.org/info/rfc4861>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<http://www.rfc-editor.org/info/rfc4944>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<http://www.rfc-editor.org/info/rfc6282>>.
- [RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <<http://www.rfc-editor.org/info/rfc6775>>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/RFC7252, June 2014, <<http://www.rfc-editor.org/info/rfc7252>>.

- [RFC7668] Nieminen, J., Savolainen, T., Isomaki, M., Patil, B., Shelby, Z., and C. Gomez, "IPv6 over BLUETOOTH(R) Low Energy", RFC 7668, DOI 10.17487/RFC7668, October 2015, <<http://www.rfc-editor.org/info/rfc7668>>.
- [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-00 (work in progress), July 2016.
- [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-01 (work in progress), July 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/>.
- [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.
- [etsi_ltn]
"ETSI Technical Committee on EMC and Radio Spectrum Matters (ERM) TG28 Low Throughput Networks (LTN)", February 2015.
- [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", Nov 2016, <URL TBD>.

[LoRaSpec1.0]

LoRa Alliance, "LoRaWAN Specification Version V1.0", Jan 2015, <<https://www.lora-alliance.org/portals/0/specs/LoRaWAN%20Specification%201R0.pdf>>.

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: Standards Track
Expires: April 26, 2017

C. Gomez
J. Paradells
UPC/i2CAT
J. Crowcroft
University of Cambridge
October 23, 2016

LPWAN Fragmentation Header
draft-gomez-lpwan-fragmentation-header-03

Abstract

LPWAN technologies are characterized by a very limited data unit and/or payload size, often one order of magnitude below the one in IEEE 802.15.4. However, many such technologies do not support layer 2 fragmentation. The 6LoWPAN fragmentation header defined in RFC 4944 represents very high overhead for LPWAN technologies, and it even does not support transporting IPv6 datagrams that require fragmentation over layer 2 technologies of a maximum payload size below 13 bytes. This specification defines an optimized fragmentation header for 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 <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 April 26, 2017.

Copyright Notice

Copyright (c) 2016 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
1.1. Conventions used in this document	3
2. FHL rules and format	3
3. Changes from RFC 4944 fragmentation header and rationale . .	5
4. IANA Considerations	6
5. Security Considerations	6
6. Acknowledgments	7
7. Annex A. Quantitative comparison of RFC 4944 fragmentation header with LFH	7
8. References	8
8.1. Normative References	8
8.2. Informative References	9
Authors' Addresses	9

1. Introduction

Low Power Wide Area Network (LPWAN) technologies are characterized, among others, by a very reduced data unit and/or payload size [I-D.minaburo-lpwan-gap-analysis]. However, many such technologies do not support layer two fragmentation, therefore the only option for these to support IPv6 (and, in particular, its MTU requirement of 1280 bytes [RFC2460]) is the use of a fragmentation mechanism at the adaptation layer below IPv6.

The 6LoWPAN fragmentation mechanism [RFC4944] is appropriate for IEEE 802.15.4-2003 (which has a frame payload size of 81 to 102 bytes). However, 6LoWPAN fragmentation it is not suitable for several LPWAN technologies. Overhead of the 6LoWPAN fragmentation header is high, considering the reduced payload size of LPWAN technologies (many of which have a maximum payload size that is one order of magnitude below that of IEEE 802.15.4-2003) and the limited energy availability of the devices using such technologies. Furthermore, the datagram offset field of the 6LoWPAN fragmentation header is expressed in increments of eight octets. The 6LoWPAN fragmentation header plus eight octets from the original datagram exceeds the available space in the layer 2 (L2) payload of some LPWAN technologies, thus 6LoWPAN fragmentation cannot be used to carry IPv6 packets over these.

This specification defines the LPWAN Fragmentation Header (LFH). While LFH has been designed for LPWAN technologies, other L2 technologies beyond the LPWAN category may benefit from using LFH.

It is expected that this specification will be used jointly with other mechanisms such as header compression.

The benefits of using LFH are the following:

-- While the 6LoWPAN fragmentation header defined in RFC 4944 has a size of 4 bytes (first fragment) or 5 bytes (subsequent fragments), LFH has a size of 2 bytes (any fragment). This reduces significantly both the L2 overhead and the adaptation layer overhead for transporting an IPv6 packet that requires fragmentation (see Annex A).

-- Because the datagram offset can be expressed in increments of a single octet, LFH enables the transport of IPv6 packets over L2 data units with a maximum payload size as small as only 3 bytes in the most extreme case.

1.1. Conventions used in this document

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

2. FHL rules and format

If an entire payload (e.g., IPv6) datagram fits within a single L2 data unit, it is unfragmented and a fragmentation header is not needed. If the datagram does not fit within a single L2 data unit, it SHALL be broken into fragments. The first fragment SHALL contain the first fragment header as defined in Figure 1.




Figure 1: First Fragment

The second and subsequent fragments (up to and including the last) SHALL contain a fragmentation header that conforms to the format shown in Figure 2.




Figure 2: Subsequent Fragments

datagram_size: This 11-bit field encodes the size of the entire IP packet before link-layer fragmentation (but after IP layer fragmentation), expressed in octets. For IPv6, the datagram size SHALL be 40 octets (the size of the uncompressed IPv6 header) more than the value of Payload Length in the IPv6 header [RFC4944] of the packet. Note that this packet may already be fragmented by hosts involved in the communication, i.e., this field needs to encode a maximum length of 1280 octets (the required by IPv6).

tag: The value of tag (datagram tag) SHALL be the same for all fragments of a payload (e.g., IPv6) datagram. The sender SHALL increment datagram_tag for successive, fragmented datagrams. The incremented value of tag SHALL wrap from 7 back to zero. This field is 3 bits long, and its initial value is not defined.

datagram_offset: This field is present only in the second and subsequent fragments and SHALL specify the offset, in increments of 1 octet, of the fragment from the beginning of the payload datagram. The first octet of the datagram (e.g., the start of the IPv6 header) has an offset of zero; the implicit value of datagram_offset in the first fragment is zero. This field is 11 bits long.

Note: the first bit of the LFH formats defined above could be used to identify an LFH header (when set to 1) or another header (when set to 0). This will need to be aligned with work-in-progress header compression specifications for LPWAN. The second bit in an LFH format determines whether a fragment is the first one or a subsequent one.

The recipient of link fragments SHALL use (1) the sender’s L2 source address (if present), (2) the destination’s L2 address (if present), (3) datagram_size, and (4) tag to identify all the fragments that belong to a given datagram.

Upon receipt of a link fragment, the recipient starts constructing the original unfragmented packet whose size is datagram_size. It uses the datagram_offset field to determine the location of the individual fragments within the original unfragmented packet. For example, it may place the data payload (except the encapsulation

header) within a payload datagram reassembly buffer at the location specified by `datagram_offset`. The size of the reassembly buffer SHALL be determined from `datagram_size`.

If a fragment recipient disassociates from its L2 network, the recipient MUST discard all link fragments of all partially reassembled payload datagrams, and fragment senders MUST discard all not yet transmitted link fragments of all partially transmitted payload (e.g., IPv6) datagrams. Similarly, when a node first receives a fragment with a given tag, it starts a reassembly timer. When this time expires, if the entire packet has not been reassembled, the existing fragments MUST be discarded and the reassembly state MUST be flushed. The reassembly timeout MUST be set to a maximum of TBD seconds).

Implementers need to be aware that in some LPWAN technologies, the MTU in use may vary over time.

3. Changes from RFC 4944 fragmentation header and rationale

This specification has used RFC 4944 fragmentation header format as a basis. The main changes introduced in this specification to the fragmentation header format defined in RFC 4944 are listed below, together with their rationale:

-- The datagram size field is only included in the first fragment. Rationale: In the RFC 4944 fragmentation header, the datagram size was included in all fragments to ease the task of reassembly at the receiver, since in an IEEE 802.15.4 mesh network, the fragment that arrives earliest to a destination is not necessarily the first fragment transmitted by the source. However, in LPWAN, such reordering effects are not expected. LPWAN technologies typically define star topology networks, peripheral to peripheral communications are not expected, and the central device is not expected to perform priority queuing operations. Nevertheless, the fragmentation format defined in this document supports limited reordering.

-- The tag size is reduced from 2 bytes to 3 bits. Rationale: Given the low bit rate, as well as the low message rate of LPWAN technologies, ambiguities due to datagram tag wrapping events are expected to occur with low probability despite the reduced tag space. The reduced tag size provides significant overhead decrease.

-- The original 1-byte RFC 4944 6LoWPAN Dispatch field is not used. Instead, two bits are used to signal an LFH header and whether a fragment is the first one or not (this, to be aligned with on-going work on header compression specifications).

-- The datagram offset size is increased from 8 bits to 11 bits.
Rationale: This allows to express the datagram offset in single-octet increments.

4. IANA Considerations

TBD

5. Security Considerations

6LoWPAN fragmentation attacks have been analyzed in the literature. Countermeasures to these have been proposed as well [HHWH].

A node can perform a buffer reservation attack by sending a first fragment to a target. Then, the receiver will reserve buffer space for the whole packet on the basis of the datagram size announced in that first fragment. Other incoming fragmented 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 fragments of multiple packets can be stored in the reassembly buffer. To further increase the attack cost, the reassembly buffer can be split into 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 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 fragment, it can send a spoofed duplicate (e.g. with random payload) to the destination. A receiver cannot distinguish legitimate from spoofed 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 fragments to be transmitted by a node, by applying content-chaining to the different fragments, based on cryptographic hash functionality. The aim of this technique is to allow a receiver to identify illegitimate fragments.

Further attacks may involve sending overlapped fragments (i.e. comprising some overlapping parts of the original datagram) or announcing a datagram size in the first fragment that does not reflect the actual amount of data carried by the fragments. Implementers should make sure that correct operation is not affected by such events.

6. Acknowledgments

In section 2, the authors have reused extensive parts of text available in section 5.3 of RFC 4944, and would like to thank the authors of RFC 4944.

The authors would like to thank Carsten Bormann, Tom Phinney, Ana Minaburo and Laurent Toutain for valuable comments that helped improve the document.

Carles Gomez has been funded in part by the Spanish Government (Ministerio de Educacion, Cultura y Deporte) through the Jose Castillejo grant CAS15/00336. 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.

7. Annex A. Quantitative comparison of RFC 4944 fragmentation header with LFH

		IPv6 datagram size (bytes)							
		11		40		100		1280	
L2 payload (bytes)		4944	LFH	4944	LFH	4944	LFH	4944	LFH
10		----	2	----	5	----	13	----	160
15		1	1	5	4	13	8	160	99
20		1	1	4	3	12	6	159	62
25		1	1	3	2	7	5	80	56
30		1	1	2	2	5	4	54	46

Figure 3: L2 overhead (in terms of L2 data units) required to transport an IPv6 datagram

		IPv6 datagram size (bytes)							
		11		40		100		1280	
L2 payload (bytes)		4944	LFH	4944	LFH	4944	LFH	4944	LFH
10		----	4	----	10	----	26	----	320
15		0	0	24	8	64	16	799	198
20		0	0	19	6	59	12	794	144
25		0	0	14	4	34	10	399	112
30		0	0	9	4	24	8	269	92

Figure 4: Adaptation layer fragmentation overhead (in bytes) required to transport an IPv6 datagram

Note 1: with the RFC 4944 fragmentation header it is not possible to transport IPv6 datagrams of the considered sizes over a 10-byte payload L2 technology.

Note 2: 11 bytes is the size of an IPv6 datagram with a 3-byte RFC 6282 compressed header (the shortest possible IPv6 header that uses global addresses), a 4-byte RFC 6282 UDP compressed header, and a CoAP message without header options and without payload.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <<http://www.rfc-editor.org/info/rfc2460>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<http://www.rfc-editor.org/info/rfc4944>>.

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

8.2. Informative References

[HHWH] Hummen et al, R., "6LoWPAN fragmentation attacks and mitigation mechanisms", 2013.

[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.

Authors' Addresses

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

Email: carlesgo@entel.upc.edu

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

Email: josep.paradells@entel.upc.edu

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

Email: jon.crowcroft@cl.cam.ac.uk

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 22, 2017

A. Minaburo, Ed.
Acklio
C. Gomez, Ed.
UPC/i2CAT
L. Toutain
Institut MINES TELECOM ; TELECOM Bretagne
J. Paradells
UPC/i2CAT
J. Crowcroft
University of Cambridge
October 19, 2016

LPWAN Survey and GAP Analysis
draft-minaburo-lpwan-gap-analysis-02

Abstract

Low Power Wide Area Networks (LPWAN) are technologies covering different applications based on long range, low bandwidth and low power operation. The use of IETF protocols in the LPWAN technologies should contribute to the deployment of a wide number of applications in an open and standard environment where actual devices using LPWAN technologies will be able to communicate. This document makes a survey of the principal characteristics of these technologies and provides the gaps for the integration on the IETF protocol stack.

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 April 22, 2017.

Copyright Notice

Copyright (c) 2016 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. Problem Statement	3
2.1. Benchmark change	4
2.2. Architecture	5
3. Analysis of gaps in current IETF protocols concerning LPWANs	6
3.1. IPv6 and LPWAN	6
3.1.1. Unicast and Multicast mapping	7
3.2. 6LoWPAN and LPWAN	7
3.2.1. 6LoWPAN Header Compression	7
3.2.2. Address Autoconfiguration	8
3.2.3. Fragmentation	8
3.2.4. Neighbor Discovery	9
3.3. 6lo and LPWAN	9
3.4. 6tisch and LPWAN	10
3.5. RoHC and LPWAN	10
3.6. ROLL and LPWAN	10
3.7. CoRE and LPWAN	11
3.8. Security and LPWAN	11
3.9. Mobility and LPWAN	11
3.9.1. NEMO and LPWAN	12
3.10. DNS and LPWAN	12
4. Annex A -- survey of LPWAN technologies	12
5. Annex B -- Security in LPWAN technologies	14
6. Acknowledgements	15
7. Normative References	15
Authors' Addresses	16

1. Introduction

LPWAN (Low-Power Wide Area Network) technologies define long range, low bit rate and low power wireless interfaces that are a kind of constrained and challenged networks [RFC7228]. They can operate in license or license-exempt bands to provide connectivity to a vast number of battery-powered devices requiring limited communications.

If the existing pilot deployments have shown the huge potential and the industrial interest in their capabilities, the loose coupling with the Internet makes the device management and network operation complex. More importantly, LPWAN devices are, as of today, with no 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 is to adapt IETF defined protocols, addressing schemes and naming spaces to this constrained environment. This document surveys the main characteristics of LPWAN technologies, and analyzes the gaps for the integration of the IETF protocol stack in the LPWAN technologies.

2. Problem Statement

The LPWANs are large-scale constrained networks in the sense of [RFC7228] with the following characteristics:

- o very small frame payload as low as 8 bytes. Typical traffic patterns are composed of a large majority of frames with payload size around 15 bytes and a small minority of up to 100 byte frames.
- o very low bandwidth, most LPWAN technologies offer a throughput between 50 bit/s to 250 kbit/s.
- o in some technologies, very limited message rate (e.g. between ~0.1 message/minute and ~1 message/minute) due to regional regulations that limit the duty cycle (e.g. from 0.1% to 10%) in some ISM bands. Some nodes will exchange less than 10 frames per day.
- o high packet loss rate, which can be the result of bad transmission conditions between nodes.
- o variable MTU for a link depending on the used L2 modulation.
- o in some cases, lack of L2 fragmentation capabilities.
- o highly asymmetric and in some cases unidirectional links.

- o ultra dense networks with thousands to tens of thousands of nodes.
- o typically, star topology networks.
- o different modulations and radio channels within the same technology.
- o sleepy nodes to preserve energy.

In the terminology of [RFC7228], these characteristics put LP-WANs into the "challenged network" category where the IP connectivity has to be redefined or modified. Therefore, LPWANs need to be considered as a separate class of networks with the following desired properties:

- o keep compatibility with current Internet:
 - * preserve the end-to-end communication principle.
 - * maintain independence from L2 technology.
 - * use or adapt protocols defined by IETF to this new environment that could be less responsive.
 - * use existing addressing spaces and naming schemes defined by IETF.
- o ensure the correspondence with the stringent LPWAN requirements, such as:
 - * limited number of messages per time unit per device.
 - * small message size, with potentially no L2 fragmentation.
 - * RTTs potentially orders of magnitude bigger than in existing constrained networks.
- o optimize the protocol stack in order to limit the number of duplicated functionalities; for instance acknowledgements should not be generated at several layers.

2.1. Benchmark change

The data transmission rate (DTR) is the amount of digital data that is moved from one device to another in a given time. In a network, the DTR can be viewed as the speed of travel of a given amount of data. The greater the bandwidth of a path the higher the DTR

(usually measured in bit per second, bit/s). For example, a low-speed connection to the Internet may have a DTR of 33.6 kilobit/s.

LPWAN DTR is lower than 1 byte/s in the most constrained links. So standard marks need to be adjusted, like for instance, instead of sending data per second we are sending the data per day. Which implies that many of the actual protocols need to be adapted to this new scale.

Timers, delays, RTTs, buffers, etc. will need to make a time shift in order to correctly perform over an LPWAN link. A concrete example could be the CoAP timers that need to be tuned properly.

2.2. Architecture


LPWAN technologies, such as LoRaWAN, NB-IOT and SIGFOX, have similar architectures but different terminology. We can identify different types of entities in a typical LPWAN network:

- o The Host, which 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 is a high density of hosts 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 AAA Server, which controls the user authentication, the applications. It is known as: Join-Server, Home Subscriber Server or Registration Authority.
- o At last we have the Application Server, known also as Packet Data Node Gateway or Network Application.

Function/ Technology	LORAWAN	NB-IOT	SIGFOX	IETF
Sensor, Actuator, device, object	End Device	User Equipment	End Point	Thing (HOST)
Transceiver Antenna	Gateway	Evolved Node B	Base Station	RADIO GATEWAY
Server	Network Server	Serving- Gateway	Service Center	Network Gateway (ROUTER)
Security Server	Join Server	Home Subscriber Server	Registration Authority	AAA SERVER
Application	Application Server	Packet Data Node Gateway	Network Application	APPLICATION SERVER

LPWAN Architecture Terminology

Figure 1



LPWAN Architecture

Figure 2

3. Analysis of gaps in current IETF protocols concerning LPWANs

3.1. IPv6 and LPWAN

IPv6 [RFC2460] 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 the LPWAN constraints. If IPv6 (with no further optimization)

were used, several LPWAN frames would be needed just to carry the header, discussion on this point is developed in the 6LoWPAN section below. Another limitation comes from the IPv6 MTU requirement, by which the layer below IP has to support packets of at least 1280 bytes [RFC2460].

IPv6 needs a configuration protocol (neighbor discovery protocol, NDP [RFC4861]) for a node to learn network parameters, and the node relation with its neighbours. This protocol generates a regular traffic with a large message size that does not fit LPWAN constraints.

3.1.1. Unicast and Multicast mapping

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.

3.2. 6LoWPAN and LPWAN

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 characterised by device constraints (in terms of processing capacity, memory, and energy availability), and specially, link constraints, such as:

- o very low 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.

3.2.1. 6LoWPAN 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 lowest 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.

3.2.2. Address Autoconfiguration

In the ambit of 6LoWPAN, traditionally, Interface Identifiers (IIDs) have been derived from link layer identifiers [RFC4944]. This allows optimisations 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.

3.2.3. 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.

3.2.4. 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 6CO, as well as the address Registration Lifetime. However, for the latter LPWANs mentioned above, 6LoWPAN Neighbor Discovery is not suitable.

3.3. 6lo and LPWAN

The 6lo WG has been reusing and adapting 6LoWPAN to enable IPv6 support over a variety of constrained node link layer technologies such as Bluetooth Low Energy (BLE), ITU-T G.9959, DECT-ULE, MS/TP-RS485, NFC or IEEE 802.11ah.

These technologies are relatively similar in several aspects to IEEE 802.15.4, which was the original 6LoWPAN target technology. 6LoWPAN has been the basis for the functionality defined by 6Lo, which has mostly used the subset of 6LoWPAN techniques most suitable for each lower layer technology, and has provided additional optimizations for technologies where the star topology is used, such as BLE 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.

3.4. 6tisch and LPWAN

The 6tisch solution is dedicated to mesh networks that operate using 802.15.4e MAC with a deterministic slotted channel. The 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.

3.5. RoHC and LPWAN

RoHC header compression mechanisms were defined for point to point multimedia channels, to reduce the header overhead of the RTP flows, it can also reduce the overhead of IPv4 or IPv6 or IPv4/v6/UDP headers. It is based on a shared context which does not require any state but packets are not routable. The context is initialised at the beginning of the communication or when it is lost. The compression is managed using a sequence number (SN) which is encoded using a window algorithm letting the reduction of the SN to 4 bits instead of 2 bytes. But this window needs to be updated each 15 packets which implies larger headers. When RoHC compression is used we talk about an average header compression size to give the performance of compression. For example, the compression start sending bigger packets than the original (52 bytes) to reduce the header up to 4 bytes (it stays here only for 15 packets, which correspond to the window size). Each time the context is lost or needs to be synchronised, packets of about 15 to 43 bytes are sent.

The RoHC header compression is not adapted to the constrained nodes of the LPWAN networks: it does not take into account the energy limitations and the transmission rate, and context is synchronised during the transmission, which does not allow a better compression.

3.6. ROLL and LPWAN

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

3.7. CoRE and LPWAN

CoRE provides a resource-oriented framework for applications intended to run on constrained IP networks. It may be necessary to adapt the protocols to take into account the duty cycling 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 acknowledgements may allow the reduction of L2 acknowledgements. 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 proves 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.

3.8. Security and LPWAN

Most of the LPWAN technologies integrate some authentication or encryption mechanisms (see Section 5) that may not have been defined by the IETF. The working group will work to integrate these mechanisms to unify management. For the technologies which are not integrating natively security protocols, it is necessary to adapt existing mechanisms to the LPWAN constraints. The AAA infrastructure brings a scalable solution. It offers a central management for the security processes, draft-garcia-dime-diameter-lorawan-00 and draft-garcia-radext-radius-lorawan-00 explain the possible security process for a LoRaWAN network. The mechanisms basically are divided in: key management protocols, encryption and integrity algorithms used. Most of the solutions do not present a key management procedure to derive specific keys for securing network and or data information. In most cases, a pre-shared key between the smart object and the communication endpoint is assumed.

3.9. Mobility and LPWAN

LPWANs 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. The mobility will imply most of the time a group of devices, which represent a network itself. The mobility concerns more the gateway than the devices.

3.9.1. NEMO and LPWAN

NEMO Mobility solutions may be used in the case where some hosts belonging to the same Network gateway will move from one point to another and that they are not aware of this mobility.

3.10. DNS and LPWAN

The purpose of the DNS is to enable applications to name things that have a global unique name. Lots of protocols are using DNS to identify the objects, especially REST and applications using CoAP. Therefore, things should be registered in DNS. DNS is probably a good topic of research for LPWAN technologies, while the matching of the name and the IP information can be used to configure the LPWAN devices.

4. Annex A -- survey of LPWAN technologies

Different technologies can be included under the LPWAN acronym. The following list is the result of a survey among the first participant to the mailing-list. It cannot be exhaustive but is representative of the current trends.

Technology	range	Throughput	MAC MTU
LoRa	2-5 km urban <15 km suburban	0.3 to 50 kbps	256 B
SIGFOX	10 km urban 50 km rural	up:100/600 bps down: 600 bps	12/ 8 B
IEEE802.15.4k LECIIM	< 20 km LoS < 5 km NoLoS	1.5 bps to 128 kbps	16/24/ 32 B
IEEE802.15.4g SUN	2-3 km LoS	4.8 kbps to 800 kbps	2047 B
RPMA	65 km LoS 20 km NoLoS	up: 624kbps down: 156kbps mob: 2kbps	64 B
DASH-7	2 km	9 kbps 55.55 kbps 166.66 kbps	256 B
Weightless-w	5 km urban	1 kbps to 10 Mbps	min 10 B
Weightless-n	<5 km urban <30 km suburban	30 kbps to 100kbps	max 20 B
Weightless-p	> 2 km urban	up to 100kbps	
NB-IoT *	<15 km	~ 200kbps	>1000B

* supports segmentation

Figure 3: Survey of LPWAN technologies

The table Figure 3 gives some key performance parameters for some candidate technologies. The maximum MTU size must be taken carefully, for instance in LoRa, it take up to 2 sec to send a 50 Byte frame using the most robust modulation. In that case the theoretical limit of 256 B will be impossible to reach.

Most of the technologies listed in the Annex A work in the ISM band and may be used for private a public networks. Weightless-W uses white spaces in the TV spectrum and NB-LTE will use licensed channels. Some technologies include encryption at layer 2.

5. Annex B -- Security in LPWAN technologies

LORAWAN

LoRaWAN provides a joining procedure called "Over the Air Activation" that enables a smart object to securely join the network, deriving the necessary keys to perform the communications securely. The messages are integrity protected and the application information is ciphered with the derived keys from the joining procedure.

The joining procedure consists of one exchange, that entails a join-request message and a join-accept message. Upon successful authentication, the smart-object and the network-server are able to derive two keys to secure the communications (AppSKey and NwkSKey)

SIGFOX

The SIGFOX radio protocol provides mechanisms to authenticate and ensure integrity of the message. This is achieved by using a unique device ID and a message authentication code, which allow ensuring that the message has been generated and sent by the device with the ID claimed in the message.

Security keys are independent for each device. These keys are associated with the device ID and they are pre-provisioned. Application data can be encrypted by the application provider.

IEEE802.15.4k and IEEE802.15.4g

There is no mention of acquiring key material to secure the communications.

DASH-7

DASH-7 defines 2 keys for specific users (root, user) and a network key. Provides network security, integrity and encryption. The process of how these keys are distributed is not explained.

RPMA

They use security algorithms and provides for mutual device authentication, message authentication and message confidentiality. No mention of how the key material is distributed.

Weightless

They offer a joining procedure to network by authenticating the smart object. Integrity of the messages, encryption and key distribution

NB-IoT

ToDo. Not Access to the specification.

6. Acknowledgements

Thanks you very much for the discussion and feedback on the LPWAN mailing list, namely, Alexander Pelov, Pascal Thubert, Samita Chakrabarti, Xavier Vilajosana, Misha Dohler, Florian Meier, Timothy J. Salo, Michael Richardson, Robert Cragie, Paul Duffy, Pat Kinney, Joaquin Cabezas and Bill Gage.

We would like also to thank Dan Garcia Carrillo and Rafael Marin Lopez for the input made for the security part.

Carles Gomez has been funded in part by the Spanish Government (Ministerio de Educacion, Cultura y Deporte) through the Jose Castillejo grant CAS15/00336. 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.

7. Normative References

- [RFC1981] McCann, J., Deering, S., and J. Mogul, "Path MTU Discovery for IP version 6", RFC 1981, DOI 10.17487/RFC1981, August 1996, <<http://www.rfc-editor.org/info/rfc1981>>.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <<http://www.rfc-editor.org/info/rfc2460>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<http://www.rfc-editor.org/info/rfc4861>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<http://www.rfc-editor.org/info/rfc4944>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<http://www.rfc-editor.org/info/rfc6282>>.

- [RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <<http://www.rfc-editor.org/info/rfc6775>>.
- [RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <<http://www.rfc-editor.org/info/rfc7228>>.
- [RFC7668] Nieminen, J., Savolainen, T., Isomaki, M., Patil, B., Shelby, Z., and C. Gomez, "IPv6 over BLUETOOTH(R) Low Energy", RFC 7668, DOI 10.17487/RFC7668, October 2015, <<http://www.rfc-editor.org/info/rfc7668>>.

Authors' Addresses

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

Email: ana@ackl.io

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

Email: carlesgo@entel.upc.edu

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

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

Email: josep.paradells@entel.upc.edu

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

Email: jon.crowcroft@cl.cam.ac.uk

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 8, 2017

A. Minaburo
Acklio
L. Toutain
Institut MINES TELECOM ; TELECOM Bretagne
July 7, 2016

RoHC applicability in LPWAN
draft-minaburo-lpwan-rohc-applicability-00

Abstract

This document makes a survey of the way to adapt the RoHC mechanisms to the LPWAN networks. It aims to show that RoHC header compression is not adapted for the constrained LPWAN networks. RoHC was defined to reduce the overhead over point-to-point connections for multimedia flows, which does not correspond to the LPWAN traffic description. RoHC is not able to reduce the use of battery and optimize the energy lifetime. If RoHC is used it will need a big effort to be adapted, there would be some problems to preserve a good transmission and packet lost will cause a context desynchronisation which implies a transmission of a complete header. LP-WAN's are different technologies covering different applications based on long range, low bandwidth and low power 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 <http://datatracker.ietf.org/drafts/current/>.

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

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

Copyright Notice

Copyright (c) 2016 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.

1. Introduction

LPWA (Low-Power Wide Area Network) technologies are a kind of constrained and challenged networks [RFC7228] based on a star topology where the device communicates directly with the gateway. The use of the Internet Protocol to communicate, specially the IPv6/UDP protocol stack creates a big overhead since the packet size of the LPWA technologies is no more than 200 bytes and in some technologies no more than 60 bytes. So the use of a header compression adapted to this very constrained networks is important. But RoHC as defined in [RFC4995] and [RFC5225] are not adapted to the constrained characteristics of the LPWAN networks. The adaptation will modify RoHC in such manner that it will not be compatible with the older versions, it will be so different that only the name will be the same. This document wants to show that it is better to think in a new solution than trying to adapt the older solution.

2. RoHC Compression

ROHC mechanism reduces the size of transmitted header by removing redundancy. ROHC mechanisms starts by classifying header fields in different classes according to their changing pattern [RFC3095]. Those fields, which are classified as inferred, are not sent, those, which are static, are sent initially and then they are not sent at all and those, which are changing, are always sent. ROHC mechanism is based on a context, which is maintained, by both ends, the compressor and the decompressor. Context keeps the entire header and ROHC's information. Each context has a context identifier (CID), which identifies the flows. ROHC defines some profiles, which define the protocol encapsulation that will be compressed. ROHC has three operation modes: Unidirectional mode (U), Optimistic mode (O) and Reliable mode (R). The U-mode is used when the link is unidirectional or when feedback is not possible. For bi-directional links, O-mode uses positive feedback packets (ACK) and R-mode use positive and negative feedback packets (ACK and NACK). ROHC always starts header compression using U-mode even if it is used in a bi-directional link. ROHC does not make retransmission when an error occurs the wrong packet is dropped. The ROHC feedback is used only

to indicate to the compressor side that there were an error and probably the context is damaged. After receiving a negative feedback compressor always reduces its compression level, which means increase the header size. ROHC compressor has three compression levels [3]: Initialization and Refresh (IR), First Order (FO) and Second Order (SO). Each compression level uses different header format packets to send the header information. In the IR compression level, it establishes the context, which contains static and dynamic header information, it is bigger than the entire header. The FO compression level gives the change pattern of dynamic fields. And, in the last compression level, SO, it sends encoded values of Sequence Number (SN) and Timestamp (TS), forming the minimal size packets. With the use of this header format packet all header fields can be generated at the other end of the link using the previously established change pattern. When some updates or errors are there, the compressor goes back to upper compression levels. It only returns to the SO compression level after retransmitting the updated information and establishing again the change pattern in the decompressor. In U-mode, the feedback channel is not used. To increase the compression level an optimistic approach is used for compressor to be sure that the context has been established at decompressor side. This means that compressor uses the same header format packet for a number of packets. As compressor does not know if context is lost it also uses two timers, to come back to FO and IR compression levels. The decompressor works at the receiving end of the link [1] and decompresses the headers based on the header fields' information of the context. Both the compressor and the decompressor use a context to store all the information about the header fields. To ensure correct decompression, the context should be always synchronized. The decompressor has three states: the first, No Context (NC), is used when there is no context synchronization, the second, Static Context (SC), is reached if the dynamic information of the context has been lost. The third, Full Context (FC), is reached when the decompressor has all the information about header fields. In FC state, the decompressor moves to the initial states as soon as it detects context damage. Decompressor uses the 'k out of n' rule by looking at the last n packets with CRC failures. If k CRC failures have occurred then it assumes context damage and transits backward to an initial state (SC or NC). The decompressor also sends feedback according to the operation mode.

2.1. Unidirectional Mode

U-mode is not only the mode where RoHC compression is initialised, it is also the only way to use RoHC in links where downlink is not available. The U-mode of RoHC is not as efficiency as the others mode of operation, because context synchronisation has to be done frequently in order to be sure that the decompressor has all the

information to reestablish the header. The decompressor state machine has three states No-Context, Static-Context and Full-Context. In the beginning the state machine initialise in No-Context state where it accepts only the IR packets, once the IR packet has arrive he goes to the Full Context state where it decompresses all the format packets. When there is an error or a packet lost and the decompressor is not able to decompress the packet he goes back to Static-context state where SO packets are dropped and only FO or IR packets are decompressed. This is because when error or lost the W-LSB for the Sequence number is lost and as in the SO header format the W-LSB SN is the only information sent, so it is not capable to find the good value. So the compressor has to refresh context as soon as it believes is needed based on error rate and parameters negotiated in the link before compression.

3. Applicability

3.1. Connection

RoHC uses a point-to-point communication with a negotiation to define the characteristics of the channel (error rate,...) and some RoHC parameters such as the size of the context identifiers, and the profiles (the header stack) will be supported for the compression in this channel. This implies to have a connection from one side to the other. The LPWAN networks do not create a channel or connection, there is not negotiation before sending any packet. Of course, RoHC accepted to do an out-of-line negotiation but in all the cases this represents a connection.

3.2. The IP address compression

In RoHC compression transmission is done in point-to-point connection then the IP addresses could be elided in the header compression processes and it is not sent in the different RoHC's packets, this implies that the packet cannot be routed or forwarded between the 2 entities.

3.3. Framework

3.3.1. Contexts

RoHC framework uses contexts, they are synchronised by the transmission information. So, sometimes the compression will send a complete header packet which also contains RoHC information doing the overhead larger than the one already generated by the IP stack.

3.3.2. Header format packets

The RoHC framework works with at least 15 different header formats making the average header size around 2 bytes. But in reality you will have larger packets. The RoHC framework add some compression information needed to identify the flow is sent in the packet as Padding (1 byte), Context ID (from 1 to 3 bytes), Profile (1 byte), RoHC format packet type (5 bits to 1 byte), CRC (from 3 bits to 1 byte) Therefore, in order to get the littlest header format the channel need to be stable and with a small error transmission, in order to keep the compressor in the SO of compression and that decompression does not send a negative acknowledgement in order to detect lost packets and reduce compression.

3.3.3. Context Synchronisation

In case of error, the decompressor will send a negative acknowledge when possible and he will drop packets until he receives a complete header. This waste of energy and compression effort is very expensive for the LPWAN node which will not be enable all the time and that expects that its information arrives without confirmation.

3.3.4. Complexity

RoHC is a very complex header compression mechanism, it was defined for wireless networks (2.5 and 3G) which use expensive radio spectrum resource and have a long RTT. The main problem solved is the use of bandwidth for multimedia applications as VoIP or VoD. in specific channels (point-to-point), where a circuit is created. The RoHC framework defines the possibility to reduce different protocols, but to be reliable, it adds a lot of 'RoHC' signalling in order to keep the context updated at both sizes (Compressor and Decompressor). To be robust it refresh information of the context regularly or when needed as the result of receiving a negative feedback. So the problem it solves is clear different from IoT. In RoHC the power and the memory are not a problem, the only important thing is to reduce the use of bandwidth that is expensive even if the node needs a lot of resources.

3.3.5. L variable or optimistic approach

RoHC uses in the Unidirectional and Optimistic mode of operation a variable to bring robustness, this approach reuses the same packet format L times in order to be sure that the Decompressor has received the information. The value of this variable is not defined, it is based on the RTT, and for 3G networks it is recommended to be 3, this means that the compressor will use 3 times each format header of each compression level before augmenting the compression rate. Then each

time the Decompressor lost the context or as periodically as timers are triggered the IR packets will be sent. Creating a big overhead for the LPWAN networks.

3.3.6. Sequence Number (SN)

RoHC use the RTP sequence number to control the missing packets and to reduce the header size. When there is no RTP header, RoHC generates a 16 bit Sequence Number to guarantee robustness. The sequence number is reduced by the W-LSB algorithm which manages the number of packets lost before the context needs to be synchronised. It also manages the number of SN bits to be sent, this corresponds to the packet format so the W-LSB can reduce compression in order to slight the window and continue working.

3.3.7. LSB Window (W-LSB)

It is a function of a reference value (v_{ref}) and 'k' (number of bits in the format packet). The LSB interpretation interval function is $f(v_{ref}, k) = [v_{ref} - p, v_{ref} + (2k - 1) - p]$; for any 'k', the 'k' LSBs must uniquely identify a value in $f(v_{ref}, k)$. 'p' is a shifting parameters that may be tuned for efficiency.

3.4. Resources Usability

When RoHC was defined there were no requirements concerning the use of battery and sporadic transmission (sleeping nodes). These two LPWAN characteristics are not taking into account in the RoHC solution. The concerning was to reduce the header overhead and to adapt the robustness of the header information to the multimedia applications where most of the time they prefer to receive erroneous information than nothing at all, that's the reason why the CRC is reduced and the UDP checksum eliminate. The concern was to keep the context synchronised to loose as less as possible the application information for reducing the use of bandwidth. RoHC is not concerned for sleepy nodes or sporadic transmission, with a very little mtu's and power optimisation.

3.5. Flow

One based characteristics of the RoHC compression is the use of flow (sequence of packets from a source to a destination), larger the flow is, better the RoHC compression performs. The RoHC mechanisms have a worse compression when there is a short flows, problems were presented for the TCP profile with shortest flows were RoHC average header size is larger than in a largest flow. The concept of flow is very important in RoHC because it takes some flow parameters to

reduce the header, like the behaviour of the: sequence number and timestamp.

3.6. Packet loss and reordering

RoHC does not support disordering in the compressed packets. The packets need to be ordered before the compression. In reality the RoHCv1 supports 'p' packets in disordered, 'p' is the parameter of the W-LSB algorithm, in practice this represents 2 packets in advance before enlarging the window in SO of compression and any in advance with a worst performance in the other levels of compression but only 1 packet late in the SO level of compression and no sequentially late packets at all in the other levels. RoHCv2 handles disordering because the mechanisms have flexible format packets where the needed information is sent. In the case of packet loss, RoHC reacts very different depending on the Mode of Operation used. In the Unidirectional mode any loss does not affect the compression but the decompressor could drop all the packets until an IR packet arrives when timer is trigger. In the Optimistic and Reliable Modes, the RoHC feedback is very helpful in order to keep robustness. When needed the decompressor will ask through a feedback the synchronisation of the context.

4. Acknowledgements

We would like also to thanks Alexander Pelov for the positive discussions

5. Annex A -- Example

ToDo

6. Normative References

- [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, <<http://www.rfc-editor.org/info/rfc3095>>.
- [RFC4995] Jonsson, L-E., Pelletier, G., and K. Sandlund, "The RObust Header Compression (ROHC) Framework", RFC 4995, DOI 10.17487/RFC4995, July 2007, <<http://www.rfc-editor.org/info/rfc4995>>.

- [RFC5225] Pelletier, G. and K. Sandlund, "RObust Header Compression Version 2 (ROHCv2): Profiles for RTP, UDP, IP, ESP and UDP-Lite", RFC 5225, DOI 10.17487/RFC5225, April 2008, <<http://www.rfc-editor.org/info/rfc5225>>.
- [RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <<http://www.rfc-editor.org/info/rfc7228>>.

Authors' Addresses

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

Email: ana@ackl.io

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

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: January 9, 2017

A. Ratilainen
Ericsson
July 8, 2016

NB-IoT characteristics
draft-ratilainen-lpwan-nb-iot-00

Abstract

Low Power Wide Area Networks (LPWAN) are wireless technologies covering different Internet of Things (IoT) applications. The common characteristics for LPWANs are large coverage, low bandwidth, small data sizes and long battery life operation. One of these technologies include Narrowband Internet of Things (NB-IoT) developed and standardized by 3GPP. This document is an informational overview to NB-IoT and gives the principal characteristics and restrictions of this technology in order to help with the IETF work for providing IPv6 connectivity to NB-IoT along with other 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 January 9, 2017.

Copyright Notice

Copyright (c) 2016 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. Overview of the NB-IoT technology	3
3. System architecture	4
4. NB-IoT worst case performance	7
5. IANA Considerations	8
6. Security Considerations	8
7. Informative References	8
Author's Address	9

1. Introduction

The purpose of this document is to provide background information and typical link characteristics about NarrowBand Internet of Things (NB-IoT) to be considered in IETF's 6LPWA work.

NB-IoT is a Low Power Wide Area (LPWA) technology being standardized by the 3GPP. NB-IoT has been developed with the following objectives in mind:

- o Improved indoor coverage
- o Support of massive number of low throughput devices
- o Low delay sensitivity
- o Ultra-low device cost
- o Low device power consumption
- o Optimized network architecture

The standardization of NB-IoT was finalized with 3GPP Release-13 in June 2016, but further enhancements for NB-IoT are worked on in the following releases, for example in the form of multicast support. 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, RLC, PDCP and RRC protocol layers respectively. The new versions of the specifications including NB-IoT are not yet available

due to novelty of the feature, but should be shortly available in the 3GPP website.

2. Overview of the NB-IoT technology

Machine type communication (MTC) refers to the emerging type of wireless communications where machine-like devices talk to each other through mobile networks or locally. Its requirements range from Massive MTC type of data with low cost, low energy consumption, small data volumes and massive numbers to critical MTC type of high reliability, very low latency and very high availability.

NB-IoT has been designed to satisfy a plethora of use cases and combination of these requirements, but especially NB-IoT targets the low-end Massive MTC scenario with following requirements: Less than 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. Highest modulation scheme is QPSK in both uplink and downlink. As the name suggests, NB-IoT uses narrowbands with the 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. On uplink multi-tone SC-FDMA is used with 15 kHz tone spacing or as a special case of SC-FDMA single tone with either 15kHz or 3.75 kHz tone spacing may be used. These schemes have been selected to reduce the User Equipment (UE) complexity.

NB-IoT can be deployed in three ways. In-band deployment means that the narrowband is multiplexed within normal LTE carrier. In Guard-band deployment the narrowband uses the unused resource blocks between two adjacent LTE carriers. Also standalone deployment is supported, where the narrowband can be located alone in dedicated spectrum, which makes it possible for example to reform the GSM carrier at 850/900 MHz for NB-IoT. All three deployment modes are meant to be used in licensed 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.

For signaling optimization, two options are introduced in addition to legacy 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 from Mobility Management Entity (MME) in core network to the

UE without interaction from base station. This means there are no Access Stratum security or header compression, as the Access Stratum is bypassed, and only limited RRC procedures.

The RRC Suspend/Resume procedures reduce the signaling overhead required for UE state transition from Idle to Connected mode in order to have a user plane transaction with the network and back to Idle state by reducing the signaling messages required compared to legacy operation

With extended DRX the RRC Connected mode DRX cycle is up to 10.24 seconds and in RRC Idle the DRX cycle can be up to 3 hours.

To recap, the following is a list of the most important characteristics of NB-IoT:

- o Narrowband operation (180 kHz bandwidth) in licensed bands, either in in-band, guard band or standalone operation mode
- o Support for 1 Data Radio Bearer (DRB) is mandatory, 2 additional DRBs are optional
- o Maximum peak rate is 60 kbps in UL and 30 kbps in DL
- o No channel access restrictions (up to 100% duty cycle)
- o The maximum size of PDCP SDU and PDCP control PDU is 1600 octets in NB-IoT
- o Data over non-access stratum is supported
- o With eDRX, DRX cycle is up to 10.24 seconds in connected mode and up to 3 hours in idle mode

3. System architecture

NB-IoT network architecture is based on the network architecture of legacy LTE, which is illustrated in Figure 1. It consists of core network, called Evolved Packet Core (EPC), Evolved UMTS Terrestrial Radio Access Network (E-UTRAN) and the User Equipment (UE). Next we take a look at the key components of EPC.




Figure 1: 3GPP network architecture

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 other 3GPP technologies.

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

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 illustration of 3GPP radio protocol architecture can be seen from Figure 2.




Figure 2: 3GPP radio protocol architecture

The radio protocol architecture of NB-IoT (and LTE) is separated into control plane and user plane. 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 EPC, passing transparently through radio network. It is responsible for authentication, security control, mobility management and bearer management.

Access Stratum (AS) is the functional layer below NAS, and in 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. 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 singlecast data. In RRC_Connected the UE has a connection to the eNodeB, the network knows the UE location on cell level and the UE may receive and transmit singlecast data. RRC_Connected mode is established, when the UE is expected to be active in the network, to transmit or receive data. Connection is released, switching to RRC_Idle, when there is no traffic to save the UE battery 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, while the UE is in RRC_Idle transparently to the eNodeB.

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 consists of 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. Lower layers in user plane are similarly RLC, MAC and physical layer performing tasks mentioned above.

4. NB-IoT worst case performance

Here we consider the worst case scenario for NB-IoT. This scenario refers to the case with high coupling loss and the UE being the least capable. The link characteristics are listed assuming such conditions.

- o 180 kHz bandwidth
- o Uplink transmission
 - * 1 Data Radio Bearer (DRB)
 - * Single-tone transmission, 3.75 kHz spacing

- o 164 dB maximum coupling loss (see Table 1

Numerology	3.75 kHz
(1) Transmit power (dBm)	23.0
(2) Thermal noise density (dBm/Hz)	-174
(3) Receiver noise figure (dB)	3
(4) Occupied channel bandwidth (Hz)	3750
(5) Effective noise power = (2) + (3) + 10*log ((4)) (dBm)	-135.3
(6) Required SINR (dB)	-5.7
(7) Receiver sensitivity = (5) + (6) (dBm)	-141.0
(8) Max coupling loss = (1) - (7) (dB)	164.0

Table 1: NB-IoT Link Budget

Under such conditions, NB-IoT may achieve data rate of roughly 200 bps.

For downlink with 164 dB coupling loss, NB-IoT may achieve higher data rates, 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 higher maximum coupling loss than 170 dB with very low bit rates.

5. IANA Considerations

This memo includes no request to IANA.

6. Security Considerations

3GPP access security is specified in [TGPP33203].

7. Informative References

[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.

[TGPP36201]

3GPP, "TS 36.201 v13.2.0 - Evolved Universal Terrestrial Radio Access (E-UTRA); LTE physical layer; General description", 2016.

[TGPP36300]

3GPP, "TS 36.300 v13.4.0 (Available soon) - Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Overall description; Stage 2", 2016.

[TGPP36321]

3GPP, "TS 36.321 v13.2.0 (Available soon) - Evolved Universal Terrestrial Radio Access (E-UTRA); Medium Access Control (MAC) protocol specification", 2016.

[TGPP36322]

3GPP, "TS 36.322 v13.2.0 (Available soon) - Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Link Control (RLC) protocol specification", 2016.

[TGPP36323]

3GPP, "TS 36.323 v13.2.0 (Available soon) - 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 (Available soon) - Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Resource Control (RRC); Protocol specification", 2016.

Author's Address

Antti Ratilainen
Ericsson
Hirsalantie 11
Jorvas 02420
Finland

Email: antti.ratilainen@ericsson.com

Network Working Group
Internet-Draft

Intended status: Informational

Expires: May 1, 2017

A. Minaburo
Acklio

L. Toutain

Institut MINES TELECOM ; TELECOM Bretagne
October 28, 2016

6LPWA Static Context Header Compression (SCHC) for CoAP
draft-toutain-lpwan-coap-static-context-hc-00

Abstract

This draft discusses the way SCHC can be applied to CoAP headers and extend the number of functions (CDF) to optimize compression.

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, 2017.

Copyright Notice

Copyright (c) 2016 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.

1. Introduction

[I-D.toutain-lpwan-ipv6-static-context-hc] defines a compression technique for LPWA network based on static context. This context is said static since the element values composing the context are not learned during packet exchanges but previously installed. The context is known by both ends. A context is composed of a set of rules (referenced by rule ids). A rule describes the header fields with some associated Target Values (TV). A Matching Operator (MO) is associated to each field. The rule is selected if all the MO matches. A Compression Decompression Function is associated to each field to define the link between the compressed and decompressed value for a specific field.

This draft discusses the way SCHC can be applied to CoAP headers and extend the number of functions (CDF) to optimize compression.

2. Compressing CoAP

CoAP [RFC7252] is an implementation of a the REST architecture for constrained devices. Gateway between CoAP and HTTP can be easily build since both protocol uses the same address space (URL), caching mechanisms and methods.

Nevertheless, if limited, the size of a CoAP header may be incompatible with LPWAN constraints and some compression may be needed to reduce the header size. CoAP compression is not straightforward. Some differences between IPv6/UDP and CoAP can be enlighten. CoAP differs from IPv6 and UDP protocols:

- o IPv6 and UDP are symmetrical protocols. The same fields are found in the request and in the answer, only location in the header may change (e.g. source and destination fields). A CoAP request is different from an answer. For instance, the URI-path option is mandatory in the request and may not be found in the response.
- o CoAP also obeys to the client/server paradigm and the compression rate can be different if the request is issued from a LPWAN node or from an external device. For instance in the former case the token size may be set to one byte. In the latter case, the token size cannot be constraint and be up to 15 byte long.
- o In IPv6, main header and UDP fields have a fixed size. In CoAP, Token size may vary from 0 to 15 bytes, length is given by a field in the header. More systematically, the options are described using the Type-Length-Value principle. Evenmore regarding the option size value, the coding will be different.

- o options type in CoAP are not defined with the same value. The Delta TLV coding makes that the type is not independant of previous option and may vary regarding the options contained in the header.

2.1. CoAP usages

A LPWAN node can either be a client or a server and sometimes both. In the client mode, the LPWAN node sends request to a server and expected answer or acknowledgements. Acknowledgements can be at 2 different levels:

- o transport level, a CON message is acknowledged by an ACK message. NON confirmable messages are not acknowledged.
- o REST level, a REST request is acknowledged by an "error" code. [RFC7967] defines an option to limit the number of acknowledgements.

Note that acknowledgement can be optimized and a REST level acknowledgement can be used as a transport level acknowledgement.

2.2. CoAP protocol analysis

CoAP defines the following fields:

- o version (2 bits): this field can be elided during a compression
- o type (2 bits): defines the type of the transport messages, 4 values are defined. Regarding the type of exchange, if only NON messages are sent or CON/ACK messages, this field can be reduced to 0 or 1 bit.
- o token length (4 bytes). The standard allows up to 15 bytes for a token length. If a fix token size is chosen, then this field can be elided. If some variation in length are needed then 1 or 2 bits could be enough for most of LPWAN applications.
- o code (8 bits). This field codes the request and the response values. CoAP represents in a more compact way, coding used in HTTP, but the coding is not optimal.
- o message id (16 bits). This value is used to acknowledge CON frames. The size of this field is computed to allow the anticipation (how many frames can be sent without acknowledgement). When a value is used, [RFC7252] defines the time before it can be reused without ambiguities. This size may

be too large for a LPWAN node sending or receiving few messages a day.

- o Token (0 to 15 bytes). Token identifies active flows. Regarding the usage (stability of in time and limited number), a short token (1 Byte) can be enough.
- o options are coded through delta-TLV. The delta-T depends of previous values, length is encoded inside the option. [RFC7252] distinguishes repeatable options that can appear several time in the header. Among them we can distinguish:
 - * list options which appear several time in the header but are exclusive such as the Accept option.
 - * cumulative options which appear several time in the header but are part of a more generic value such as Uri-Path and Uri-Query.

For a given flow some value options are stable through time. Observe, ETag, If-Match, If-None-Match and Size varies in each message. Options can be stored in a SCHC context and cumulative options can be stored globally.

The CoAP protocol must not be altered by the compression/decompression phase, but if no semantic is attributed to a value, it may be changed during this phase. For instance the compression phase may reduce the size of a token but must maintain its unicity. The decompressor will not be able to restore the original value but behavior will remain the same. If no special semantic is assigned to the token, this will be transparent. If a special semantic is assigned to the token, its compression may not be possible.

This implies that the compressor/decompressor must be aware of the protocol state machine and do not processes request and response the same way.

A conservative compression leaves the field value unchanged. Non conservative compression can be used when a CoAP implementation has not been defined to work specifically with LPWAN network and uses large values for fields.

2.2.1. CoAP Compression Decompression Function

To compress more efficiently CoAP message, several Compression/Decompression Function (CDF) must be defined.

2.2.1.1. Static-mapping

The goal of static-mapping is to reduce the size of a field by allocating shorter value. The mapping is known by both ends and stored in a table in both end context. The Static-mapping is conservative.

Static-mapping may be applied to several fields. For instance the type field may be reduced from 2 bits to 1 bit if only CON/ACK type is used, but the main benefit is compressing the code field.

The CoAP code field defines a tricky way to ensure compatibility with HTTP values. Nevertheless only 21 values are defined by [RFC7252] compared to the 255 possible values. So it could efficiently be coded on 5 bits. To allow flexibility and evolution if new codes are introduced, a static mapping table is associated to each instance of this function.

Figure 1 gives a possible mapping, it can be changed to add new codes or reduced if some values are never used by both ends.

Code	Description	Mapping
0.00		0x00
0.01	GET	0x01
0.02	POST	0x02
0.03	PUT	0x03
0.04	DELETE	0x04
0.05	FETCH	0x05
0.06	PATCH	0x06
0.07	iPATCH	0x07
2.01	Created	0x08
2.02	Deleted	0x09
2.03	Valid	0x0A
2.04	Changed	0x0B
2.05	Content	0x0C
4.00	Bad Request	0x0D
4.01	Unauthorized	0x0E
4.02	Bad Option	0x0F
4.03	Forbidden	0x10
4.04	Not Found	0x11
4.05	Method Not Allowed	0x12
4.06	Not Acceptable	0x13
4.12	Precondition Failed	0x14
4.13	Request Entity Too Large	0x15
4.15	Unsupported Content-Format	0x16
5.00	Internal Server Error	0x17
5.01	Not Implemented	0x18
5.02	Bad Gateway	0x19
5.03	Service Unavailable	0x1A
5.04	Gateway Timeout	0x1B
5.05	Proxying Not Supported	0x1C

Figure 1: CoAP code mapping

This CDF can also be applied to path to send a reference instead of the path value.

2.2.1.2. Remapping

With dynamic mapping, the mapping is done dynamically, which means that the other end has no way to learn the original value. This function is not conservative. The mapping context must be stored in a reliable way on the compressor, if lost the session with LPWAN node

will be lost, which can generate a traffic increase on the LPWA network.

This function converts a large number to a smaller one and maintain bi-directional mapping. If the field has no semantic, such as a CoAP token or a message ID, this will reduce the size of the information sent on the link. This mapping only applies for request compression, answers must keep the value original value.

For instance a compression receives a CoAP request with a large token. The compressor reduces the token size by allocating a unused value in a smaller space. When the response come back, the compressor exchange the smallest token with the original one.

This mean that the compressor must be aware of the CoAP state machine, to identify a request and its associated response, but also determine when a token value can be reused.

2.2.1.3. Reduce-entropy

Reduce-entropy is a non-conservative function. the goal is to minimize the increase in a field value. It may be used for the observe option, all increase in the original sequence number will lead to an increase of 1 in the compressed value.

For instance a LPWAN node is a CoAP server and receives Observe responses coming from an outside client. The client uses a clock to generate Observe sequence number. If that value has non particular meaning for the CoAP server, increase of 1 will not change the protocol behavior. Reordering works the same way as for original Observe.

2.2.2. CoAP mandatory header

Figure 2 proposes some function assignments to the CoAP header fields.

Field	Function	Behavior
version	not-sent	version is always the same
type	value-sent	if all the types are used
sed	static-mapping	to reduce to one bit if 2 type are used
	not-sent	if only one type is used (e.g. NON)
token length	not-sent	no tokens or fixed size
	compute-token-length	if token size is reduced
	value-sent	token is sent integrally
code	value-sent	no modification
	static-mapping	code size reduction
message id	value-sent	no modification
token	remapping	reduces message id size
Content-Format	value-sent	no modification
Accept	not-sent	defined in the rule
Max-Age	static-mapping	map the possible value
Path:	value-sent	no modification
Uri-Host+Uri-Port+	not-sent	defined in the rule
Uri-Path*+Uri-Query*	static-mapping	a value to define a path
Proxy-Uri		Note: only the full path is stored in
Proxy-Scheme		context
ETag	value-sent	Always sent
Location-Path		




Figure 2: SCHC functions' example assignment for CoAP

2.2.3. Examples of CoAP header compression

2.2.3.1. 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 ES. For this simple scenario, the rules are described Figure 3

```
rule id 1
+-----+-----+-----+-----+-----+
| Field      | TV      | MO      | CDF      | Sent      |
+-----+-----+-----+-----+-----+
| CoAP version | 01      | =        | not-sent  |           | |
| CoAP Type    |         |          | value-sent| TT        |
| CoAP TKL     | 0000    | =        | not-sent  |           |
| CoAP Code    |         |          | static-map| CC CCC    |
| CoAP MID     |         |          | dynamic-map|           | M-ID      |
| CoAP Path    | /path   |          | not-sent  |           |
+-----+-----+-----+-----+-----+
```

Figure 3: CoAP Context to compress header without token

Figure 3 gives a simple compression rule for CoAP headers without tokens.

The version fields and Token Length are elided. Code has shrunk to 5 bits using the static-mapping function. Message-ID has shrunk to 9 bits to preserve alignment on byte boundary.

Figure 4 shows the time diagram of the exchange. A LPWAN Application Server sends a CON message. Compression reduces the header sending only the Type, a mapped code and the Message ID is change to a value on 9 bits. The receiver decompress the header. The message ID value is changed.

The CON message is a request, therefore the LC process to a dynamic mapping. When the ES receives the ACK message, this will not initiate locally a the message ID mapping since it is a response. The LC receives the ACK and uncompress it to restore the original value. Dynamic Mapping context lifetime follows the same rules as message ID duration.




Figure 4: Compression with global addresses

Note that the compressor and decompressor must understand the CoAP protocol:

- o The LC compressor detects a new transport request and allocate a new dynamic mapping value.
- o When receiving a response the ES compressor ES detects that this is a response (type=2) therefore the message ID value in unchanged.
- o The upstream compressor detects that is an REST answer (code 2.05) therefore the path option is not inserted in the uncompress header

2.2.3.2. Exchange with token

The following scenario introduces tokens. The LC manages two remapping contexts. One for Message ID and the other for token. ES manages one context for Message ID. Mapping is triggered by the reception of CON messages to compress or CoAP requests to compress. Note that the compressed message ID size has been reduced to 7 bits, compared to the previous example, to maintain byte boundary alignment.

Field	Function	Ctxt Value	Sent compressed
CoAP version	not-sent		
CoAP Type	value-sent		TT
CoAP TKL	compute-token-length		LL
CoAP Code	map-code	mapping table	CCCC C
CoAP MID	remapping	7 bits	M-ID
CoAP Token	remapping	8 bits	token
CoAP Path	not-sent	/data/humidity	

Figure 5: CoAP Context to compress header with token




Figure 6: Compression with token

3. Normative References

[I-D.toutain-lpwan-ipv6-static-context-hc]
 Minaburo, A. and L. Toutain, "LPWAN Static Context Header Compression (SCHC) for IPv6 and UDP", draft-toutain-lpwan-ipv6-static-context-hc-00 (work in progress), September 2016.

- [RFC1332] McGregor, G., "The PPP Internet Protocol Control Protocol (IPCP)", RFC 1332, DOI 10.17487/RFC1332, May 1992, <<http://www.rfc-editor.org/info/rfc1332>>.
- [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, <<http://www.rfc-editor.org/info/rfc3095>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<http://www.rfc-editor.org/info/rfc4944>>.
- [RFC4997] Finking, R. and G. Pelletier, "Formal Notation for RObust Header Compression (ROHC-FN)", RFC 4997, DOI 10.17487/RFC4997, July 2007, <<http://www.rfc-editor.org/info/rfc4997>>.
- [RFC5225] Pelletier, G. and K. Sandlund, "RObust Header Compression Version 2 (ROHCv2): Profiles for RTP, UDP, IP, ESP and UDP-Lite", RFC 5225, DOI 10.17487/RFC5225, April 2008, <<http://www.rfc-editor.org/info/rfc5225>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <<http://www.rfc-editor.org/info/rfc6282>>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/RFC7252, June 2014, <<http://www.rfc-editor.org/info/rfc7252>>.
- [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, <<http://www.rfc-editor.org/info/rfc7967>>.

Authors' Addresses

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

Email: ana@ackl.io

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

Network Working Group
Internet-Draft
Intended status: Informational
Expires: March 27, 2017

A. Minaburo
Acklio
L. Toutain
Institut MINES TELECOM ; TELECOM Bretagne
September 23, 2016

LPWAN Static Context Header Compression (SCHC) for IPv6 and UDP
draft-toutain-lpwan-ipv6-static-context-hc-00

Abstract

This document describes a header compression scheme for IPv6, IPv6/UDP based on static contexts. This technique is especially tailored for LPWA networks and could be extended to other protocol stacks.

During the IETF history several compression mechanisms have been proposed. First mechanisms, such as RoHC, are using a context to store header field values and send smaller incremental differences on the link. Values in the context evolve dynamically with information contained in the compressed header. The challenge is to maintain sender's and receiver's contexts synchronized even with packet losses. Based on the fact that IPv6 contains only static fields, 6LoWPAN developed an efficient context-free compression mechanisms, allowing better flexibility and performance.

The Static Context Header Compression (SCHC) combines the advantages of RoHC context which offers a great level of flexibility in the processing of fields, and 6LoWPAN behavior to elide fields that are known from the other side. Static context means that values in the context field do not change during the transmission, avoiding complex resynchronization mechanisms, incompatible with LPWA characteristics. In most of the cases, IPv6/UDP headers are reduced to a small identifier.

This document focuses on IPv6/UDP headers compression, but the mechanism can be applied to other protocols such as CoAP. It will be described in a separate document.

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 March 27, 2017.

Copyright Notice

Copyright (c) 2016 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.

1. Introduction

Headers compression is mandatory to bring the internet protocols to the node within a LPWA network [I-D.minaburo-lp-wan-gap-analysis].

Nevertheless, LPWA networks offer good properties for an efficient header compression:

- o Topology is star oriented, therefore all the packets follows the same path. For the needs of this draft, the architecture can be summarized to End-Systems (ES) exchanging information with LPWAN Application Server (LA). The exchange goes through a single LPWA Compressor (LC). In most of the cases, End Systems and LC form a star topology. ESs and LC maintain a static context for compression. Static context means that context information is not learned during the exchange.
- o Traffic flows are mostly deterministic, since End-Systems embed built-in applications. Contrary to computers or smartphones, new applications cannot be easily installed.

First mechanisms such as RoHC use a context to store header field values and send smaller incremental differences on the link. The first version of RoHC targeted IP/UDP/RTP stack. RoHCv2 extends the principle to any protocol and introduces a formal notation [RFC4997] describing the header and associating compression functions to each

field. To be efficient the sender and the receiver must check that the context remains synchronized (i.e. contains the same values). Context synchronization imposes to periodically send a full header or at least dynamic fields. If fully compressed, the header can be compatible with LPWA constraints. However, the first exchanges or context resynchronisations impose to send uncompressed headers, which may be bigger than the original one. This will force the use of inefficient fragmentation mechanisms. For some LPWA technologies, duty cycle limits can also delay the resynchronization. Figure 1 illustrates this behavior.




Figure 1: RoHC Compressed Header size evolution.

On the other hand, 6LoWPAN [RFC4944] is context-free based on the fact that IPv6, its extensions or UDP headers do not contain incremental fields. The compression mechanism described in [RFC6282] is based on sending a 2-byte bitmap, which describes how the header should be decompressed, either using some standard values or sending information after this bitmap. [RFC6282] also allows for UDP compression.

In the best case, when Hop limit is a standard value, flow label, DiffServ fields are set to 0 and Link Local addresses are used over a single hop network, the 6LoWPAN compressed header is reduced to 4 bytes. This compression ratio is possible because the IID are derived from the MAC addresses and the link local prefix is known from both sides. In that case, the IPv6 compression is 4 bytes and UDP compression is 2 bytes, which fills half of the payload of a SIGFOX frame, or more than 10% of a LoRaWAN payload (with spreading factor 12).

The Static Context Header Compression (SCHC) combines the advantages of RoHC context, which offers a great level of flexibility in the

processing of fields, and 6LoWPAN behavior to elide fields that are known from the other side. Static context means that values in the context field do not change during the transmission, avoiding complex resynchronization mechanisms, incompatible with LPWA characteristics. In most of the cases, IPv6/UDP headers are reduced to a small context identifier.

2. Static Context Header Compression

Static Context Header Compression (SCHC) avoids context synchronization, which is the most bandwidth-consuming operation in RoHC. Based on the fact that the nature of data flows is highly predictable in LPWA networks, a static context may be stored on the End-System (ES). The other end, the LPWA Compressor (LC) can learn the context through a provisioning protocol during the identification phase (for instance, as it learns the encryption key).

The context contains a list of rules (cf. Figure 2). Each rule contains itself a list of field descriptions composed of a target value (TV), a matching operator (MO) and a Compression/Decompression Function (CDF).




Figure 2: Compression Decompression Context

The rule does not describe the compressed/decompressed packet format which must be known from the compressor/decompressor. The rule just describes the compression/decompression behavior for a field.

The main idea of the compression scheme is to send the rule number (or rule id) to the other end instead of known field values.

Matching a field with a value and header compression are related operations; If a field matches a rule containing the value, it is not necessary to send it on the link. Since contexts are synchronized, reading the rule's value is enough to reconstruct the field's value at the other end.

On some other cases, the value need to be sent on the link to inform the other end. The field value may vary from one packet to another, therefore the field cannot be used to select the rule id.

2.1. Simple Example

A simple header is composed of 3 fields (F1, F2, F3). The compressor receives a packet containing respectively [F1:0x00, F2:0x1230, F3:0xABC0] in those fields. The Matching Operators (as defined in Section 3) allow to select Rule 5 as represented in Figure 3; F1 value is ignored and F2 and F3 packet field values are matched with those stored in the rule Target Values.

Rule 5			
	Target Value	Matching Operator	Comp/Decomp Fct
F1	0x00	Ignore	not-sent
F2	0x1230	Equal	not-sent
F3	0xABC0	Equal	not-sent

Figure 3: Matching Rule

The Compression/Decompression Function (as defined in Section 4) describes how the fields are compressed. In this example, all the fields are elided and only the rule number has to be sent to the other end.

The decompressor receives the rule number and reconstructs the header using the values stored in the Target Value column.

Note that F1 value will be set to 0x00 by the decompressor, even if the original header field was carrying a different value.

To allow a range of values for field F2 and F3, the MSB() Matching Operator and LSB() Compression/Decompression Function can be used (as defined in Section 3 and Section 4). In that case the rule will be rewritten as defined in Figure 4.

Rule 5			
	Target Value	Matching Operator	Comp/Decomp Fct
F1	0x00	Ignore	not-sent
F2	0x1230	MSB(12)	LSB(4)
F3	0xABC0	MSB(12)	LSB(4)

Figure 4: Matching Rule

In that case, if a packet with the following header fields [F1:0x00, F2:0x1234, F3:0xABCD] arrives to the compressor, the new rule 5 will be selected and sent to the other end. The compressed header will be composed of the single byte [0x4D]. The decompressor receives the compressed header and follows the rule to reconstruct [0x00, 0x1234, 0xABCD] applying a OR operator between the target value stored in the rule and the compressed field value sent.

2.2. Packet processing

The compression/decompression process follows several steps:

- o compression rule selection: the goal is to identify which rule will be used to compress the headers. To each field is associated a matching rule for compression. Each header field's value is compared to the corresponding target value stored in the rule for that field using the matching operator. If all the fields satisfied the matching operator, the packet is processed using this Compression Decompression Function functions. Otherwise the next rule is tested. If no eligible rule is found, then the packet is dropped.
- o sending: The rule number is sent to the other end followed by data resulting from the field compression. The way the rule number is sent depends of the layer two technology and will be specified in a specific document. For exemple, it can either be included in a Layer 2 header or sent in the first byte of the L2 payload.
- o decompression: The receiver identifies the sender through its device-id (e.g. MAC address) and select the appropriate rule through the rule number. It applies the compression decompression function to reconstruct the original header fields.

3. Matching operators

It may exist some intermediary cases, where part of the value may be used to select a field and a variable part has to be sent on the link. This is true for Least Significant Bits (LSB) where the most significant bit can be used to select a rule id and the least significant bits have to be sent on the link.

Several matching operators are defined:

- o equal: a field value in a packet matches with a field value in a rule if they are equal.
- o ignore: no check is done between a field value in a packet and a field value in the rule. The result is always true.
- o MSB(length): a field value of length T in a packet matches with a field value in a rule if the most significant "length" bits are equal.

4. Compression Decompression Functions (CDF)

The Compression Decompression Functions (CDF) describe the action taken during the compression and inversely the action taken by the decompressor to restore the original value.

Function	Compression	Decompression
not-sent	elided	use value stored in ctxt
value-sent	send	build from received value
LSB(length)	send LSB	ctxt value OR rcvd value
compute-IPv6-length	elided	compute IPv6 length
compute-UDP-length	elided	compute UDP length
compute-UDP-checksum	elided	compute UDP checksum
ESiid-DID	elided	build IID from L2 ES addr
LAiid-DID	elided	build IID from L2 LA addr

Figure 5: Compression and Decompression Functions

Figure 5 lists all the functions defined to compress and decompress a field. The first column gives the function's name. The second and third columns outlines the compression/decompression process.

As with 6LoWPAN, the compression process may produce some data, where fields that were not compressed (or were partially compressed) will be sent in the order of the original packet. Information added by the compression phase must be aligned on byte boundaries, but each individual compression function may generate any size.

Field	Comp Decomp Fct	Behavior
IPv6 version IPv6 DiffServ IPv6 FL IPv6 NH	not-sent value-sent	The value is not sent, but each end agrees on a value, which can be different from 0. Depending on the matching operator, the entire field value is sent or an adjustment to the context value
IPv6 Length	compute-IPv6-length	Dedicated fct to reconstruct value
IPv6 Hop Limit	not-sent+MO=ignore	The receiver takes the value stored in the context. It may be different from one originally sent, but in a star topology, there is no risk of loops
	not-sent+matching	Receiver and sender agree on a specific value.
	value-sent	Explicitly sent
IPv6 ESPrefix IPv6 LAPrefix	not-sent value-sent	The 64 bit prefix is stored on the context Explicitly send 64 bits on the link
IPv6 ESiid IPv6 LAiid	not-sent	IID is not sent, but stored in the context
	ESiid-DID LAiid-DID value-sent	IID is built from the ES/LA Dev. ID IID is explicitly sent on the link. Size depends of the L2 technology
UDP ESport UDP LApport	not-sent value-sent LSB(length)	In the context Send the 2 bytes of the port number or least significant bits if MSB matching is specified in the matching operator.
UDP length	compute-UDP-length	Dedicated fct to reconstruct value
UDP Checksum	compute-UDP-checksum	Dedicated fct to reconstruct value

Figure 6: SCHC functions' example assignment for IPv6 and UDP

Figure 6 gives an example of function assignment to IPv6/UDP fields.

4.1. Compression Decompression Functions (CDF)

4.1.1. not-sent

The compressor do not sent the field value on the link. The decompressor restore the field value with the one stored in the matched rule.

4.1.2. value-sent

The compressor send the field value on the link, if the matching operator is "=". Otherwise the matching operator indicates the information that will be sent on the link. For a LSB operator only the Least Significant Bits are sent.

4.1.3. LSB(length)

The compressor sends the "length" Least Significant Bits. The decompressor combines with a OR operator the value received with the Target Value.

4.1.4. ESiid-DID, LAiid-DID

These functions are used to process respectively the End System and the LA Device Identifier (DID). The IID value is computed from device ID present in the Layer 2 header. The computation depends on the technology and the device ID size.

4.1.5. Compute-*

These functions compute the field value based on received information. They are elided during the compression and reconstructed during the decompression.

- o compute-ipv6-length: compute the IPv6 length field as described in [RFC2460].
- o compute-udp-length: compute the IPv6 length field as described in [RFC0768].
- o compute-udp-checksum: compute the IPv6 length field as described in [RFC0768].

5. Examples

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

5.1. IPv6/UDP compression in a star topology

The most common case will be a LPWA end-system embeds some applications running over CoAP. In this example, the first flow is for the device management based on CoAP using Link Local addresses and UDP ports 123 and 124. The second flow will be a CoAP server for measurements done by the end-system (using ports 5683) and Global Addresses alpha::IID/64 to beta::1/64. The last flow is for legacy applications using different ports numbers, the destination is gamma::1/64.

Figure 7 presents the protocol stack for this end-system. IPv6 and UDP are represented with dotted lines since these protocols are compressed on the radio link.




Figure 7: Simplified Protocol Stack for LP-WAN

Note that in some LPWA technologies, only End Systems have a device ID . Therefore it is necessary to define statically an IID for the Link Local address for the LPWA Compressor.

Rule 0

Field	Value	Match	Function	Sent
IPv6 version	6	equal	not-sent	
IPv6 DiffServ	0	equal	not-sent	
IPv6 Flow Label	0	equal	not-sent	
IPv6 Length		ignore	comp-IPv6-1	
IPv6 Next Header	17	equal	not-sent	
IPv6 Hop Limit	255	ignore	not-sent	
IPv6 ESprefix	FE80::/64	equal	not-sent	
IPv6 ESiid		ignore	ESiid-DID	
IPv6 LCprefix	FE80::/64	equal	not-sent	

IPv6 LAiid	:::1	equal	not-sent		
UDP ESport	123	equal	not-sent		
UDP LAport	124	equal	not-sent		
UDP Length		ignore	comp-UDP-l		
UDP checksum		ignore	comp-UDP-c		

Rule 1

Field	Value	Match	Function	Sent
IPv6 version	6	equal	not-sent	
IPv6 DiffServ	0	equal	not-sent	
IPv6 Flow Label	0	equal	not-sent	
IPv6 Length		ignore	comp-IPv6-l	
IPv6 Next Header	17	equal	not-sent	
IPv6 Hop Limit	255	ignore	not-sent	
IPv6 ESprefix	alpha/64	equal	not-sent	
IPv6 ESiid		ignore	ESiid-DID	
IPv6 LAprefix	beta/64	equal	not-sent	
IPv6 LAiid	:::1000	equal	not-sent	
UDP ESport	5683	equal	not-sent	
UDP LAport	5683	equal	not-sent	
UDP Length		ignore	comp-UDP-l	
UDP checksum		ignore	comp-UDP-c	

Rule 2

Field	Value	Match	Function	Sent
IPv6 version	6	equal	not-sent	
IPv6 DiffServ	0	equal	not-sent	
IPv6 Flow Label	0	equal	not-sent	
IPv6 Length		ignore	comp-IPv6-l	
IPv6 Next Header	17	equal	not-sent	
IPv6 Hop Limit	255	ignore	not-sent	
IPv6 ESprefix	alpha/64	equal	not-sent	
IPv6 ESiid		ignore	ESiid-DID	
IPv6 LAprefix	gamma/64	equal	not-sent	
IPv6 LAiid	:::1000	equal	not-sent	
UDP ESport	8720	MSB(12)	LSB(4)	lsb
UDP LAport	8720	MSB(12)	LSB(4)	lsb
UDP Length		ignore	comp-UDP-l	
UDP checksum		ignore	comp-UDP-c	

+=====+=====+=====+=====+=====+

Figure 8: Context rules

All the fields described in the three rules Figure 8 are present in the IPv6 and UDP headers. The ESDevice-ID value is found in the L2 header.

The second and third rules use global addresses. The way the ES learn the prefix is not in the scope of the document. One possible way is to use a management protocol to set up in both end rules the prefix used on the LPWA network.

The third rule compresses port numbers on 4 bits. This value is selected to maintain alignment on byte boundaries for the compressed header.

6. Acknowledgements

Thanks to Dominique Barthel, Arunprabhu Kandasamy, Antony Markovski, Alexander Pelov, Juan Carlos Zuniga for useful design consideration.

7. Normative References

- [I-D.minaburo-lp-wan-gap-analysis]
Minaburo, A., Pelov, A., and L. Toutain, "LP-WAN GAP Analysis", draft-minaburo-lp-wan-gap-analysis-01 (work in progress), February 2016.
- [RFC0768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, DOI 10.17487/RFC0768, August 1980, <<http://www.rfc-editor.org/info/rfc768>>.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <<http://www.rfc-editor.org/info/rfc2460>>.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <<http://www.rfc-editor.org/info/rfc4944>>.
- [RFC4997] Finking, R. and G. Pelletier, "Formal Notation for RObust Header Compression (ROHC-FN)", RFC 4997, DOI 10.17487/RFC4997, July 2007, <<http://www.rfc-editor.org/info/rfc4997>>.

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

Authors' Addresses

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

Email: ana@ackl.io

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

LPWAN Working Group
Internet-Draft
Intended status: Informational
Expires: June 7, 2018

JC. Zuniga
B. Ponsard
SIGFOX
December 04, 2017

SIGFOX System Description
draft-zuniga-lpwan-sigfox-system-description-04

Abstract

This document presents an overview of the network architecture and system characteristics of a typical SIGFOX Low Power Wide Area Network (LPWAN). It is intended to be used as background information by the IETF LPWAN group when defining system requirements of different LPWAN technologies that are suitable to support common IP services.

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 June 7, 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 (<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. System Architecture	3
4. Radio Spectrum	5
5. Radio Protocol	5
5.1. Uplink	6
5.1.1. Uplink Physical Layer	6
5.1.2. Uplink MAC Layer	6
5.2. Downlink	7
5.2.1. Downlink Physical Layer	7
5.2.2. Downlink MAC Layer	7
5.3. Synchronization between Uplink and Downlink	8
6. Network Deployment	8
7. IANA Considerations	9
8. Security Considerations	9
9. Acknowledgments	9
10. Informative References	9
Authors' Addresses	10

1. Introduction

This document presents an overview of the network architecture and system characteristics of a typical SIGFOX LPWAN, which is in line with the terminology and specifications defined by ETSI [etsi_unb]. It is intended to be used as background information by the IETF LPWAN group when defining system requirements of different LPWANs that are suitable to support common IP services.

LPWAN technologies are a subset of IoT systems which specifically enable long range data transport (e.g. distances up to 50 km in open field), are capable to communicate with underground equipment, and require minimal power consumption. Low throughput transmissions combined with advanced signal processing techniques provide highly effective protection against interference. LPWAN technologies can also cooperate with cellular networks to address use cases where redundancy, complementary or alternative connectivity is needed.

Because of these characteristics, LPWAN systems are particularly well adapted for low throughput IoT traffic. SIGFOX LPWAN autonomous battery-operated devices send only a few bytes per day, week or month in an asynchronous manner and without the needed for central coordination, which allows them to remain on a single battery for up to 10-15 years.

2. Terminology

The following terms are used throughout this document:

Base Station (BS) - A Base Station is a radio hub, relaying messages between DEVs and the SC.

Device Application (DA) - An application running on the DEV or EP.

Device (DEV) - A Device (aka end-point) is a leaf node of a LPWAN that communicates application data between the local device application and the network application.

End Point (EP) - An End Point (aka device) is a leaf node of a LPWAN that communicates application data between the local device application and the network application.

Low-Power Wide-Area Network (LPWAN) - A system comprising several BSs and millions/billions of DEVs, characterized by the extreme low-power consumption, long-range of transmission, and typically connected in a star network topology.

Network Application (NA) - An application running in the network and communicating with the device(s).

Registration Authority (RA) - The Registration Authority is a central entity that contains all allocated and authorized Device IDs.

Service Center (SC) - The SIGFOX network has a single service centre. The SC performs the following functions:

- * DEVs and BSs management
- * DEV authentication
- * Application data packets forwarding
- * Cooperative reception support

3. System Architecture

Figure 1 depicts the different elements of the system architecture:




Figure 1: 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 Stations (BSs) and the Internet, as well as the control and management of the BSs and Devices. The RA is in charge of the Device 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 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 without needing to signal for handover or roaming.

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.

4. Radio Spectrum

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].

At present, 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.

5. Radio Protocol

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.

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.

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.

5.1. Uplink

5.1.1. Uplink Physical Layer

The main radio characteristics of the UNB uplink transmission are:

- o Occupied bandwidth: 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 maximum mean transmission time of 1%.

5.1.2. Uplink MAC Layer

The format of the uplink frame is the following:




Figure 2: 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)

5.2. Downlink

5.2.1. Downlink Physical Layer

The main radio characteristics of the UNB downlink transmission are:

- o Occupied bandwidth: 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: Centre frequency of downlink transmission are 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.

5.2.2. Downlink MAC Layer

The format of the downlink frame is the following:




Figure 3: 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)

5.3. Synchronization between Uplink and Downlink

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 the ISM bands. Under the strictest regulations, the system can allow a maximum of 140 uplink messages and 4 downlink messages per device. These restrictions can be slightly relaxed depending on system conditions and the specific regulatory domain of operation.

6. Network Deployment

As of today, SIGFOX's network has been fully deployed in 17 countries, with ongoing deployments on 29 other countries, giving in total a geography of 2.6 million square kilometers, containing 660 million people. The single core network model allows devices to connect in any country, without any notion of roaming or handover.

The vast majority of the current applications are sensor-based, requiring solely uplink communications, followed by actuator-based applications, which make use of bidirectional (i.e. uplink and downlink) communications.

Similar to other LPWAN technologies, the sectors that currently benefit from the low-cost, low-maintenance and long battery life are agricultural and environment, public sector (smart cities, education, security, etc.), industry, utilities, retail, home and lifestyle, health and automotive.

7. IANA Considerations

N/A.

8. Security Considerations

Due to the nature of low-complexity devices, 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.

9. Acknowledgments

The authors would like to thank Olivier Peyrusse for the useful inputs and discussions about ETSI UNB SRD.

10. Informative References

[arib_ref]

"ARIB STD-T108 (Version 1.0): 920MHz-Band Telemeter, Telecontrol and data transmission radio equipment.", February 2012.

[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.

[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.

[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.

Authors' Addresses

Juan Carlos Zuniga
SIGFOX
425 rue Jean Rostand
Labege 31670
France

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

Benoit Ponsard
SIGFOX
425 rue Jean Rostand
Labege 31670
France

Email: Benoit.Ponsard@sigfox.com
URI: <http://www.sigfox.com/>