

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45

LoRaWAN Specification

Copyright © 2016 LoRa Alliance, Inc. All rights reserved.

NOTICE OF USE AND DISCLOSURE

Copyright © LoRa Alliance, Inc. (2015, 2016). All Rights Reserved.

The information within this document is the property of the LoRa Alliance (“The Alliance”) and its use and disclosure are subject to LoRa Alliance Corporate Bylaws, Intellectual Property Rights (IPR) Policy and Membership Agreements.

Elements of LoRa Alliance specifications may be subject to third party intellectual property rights, including without limitation, patent, copyright or trademark rights (such a third party may or may not be a member of LoRa Alliance). The Alliance is not responsible and shall not be held responsible in any manner for identifying or failing to identify any or all such third party intellectual property rights.

This document and the information contained herein are provided on an “AS IS” basis and THE ALLIANCE DISCLAIMS ALL WARRANTIES EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO (A) ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OF THIRD PARTIES (INCLUDING WITHOUT LIMITATION ANY INTELLECTUAL PROPERTY RIGHTS INCLUDING PATENT, COPYRIGHT OR TRADEMARK RIGHTS) OR (B) ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE OR NON-INFRINGEMENT.

IN NO EVENT WILL THE ALLIANCE BE LIABLE FOR ANY LOSS OF PROFITS, LOSS OF BUSINESS, LOSS OF USE OF DATA, INTERRUPTION OF BUSINESS, OR FOR ANY OTHER DIRECT, INDIRECT, SPECIAL OR EXEMPLARY, INCIDENTAL, PUNITIVE OR CONSEQUENTIAL DAMAGES OF ANY KIND, IN CONTRACT OR IN TORT, IN CONNECTION WITH THIS DOCUMENT OR THE INFORMATION CONTAINED HEREIN, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

The above notice and this paragraph must be included on all copies of this document that are made.

LoRa Alliance, Inc.
2400 Camino Ramon, Suite 375
San Ramon, CA 94583

Note: All Company, brand and product names may be trademarks that are the sole property of their respective owners.



LoRaWAN™ Specification

Authors:

N. Sornin (Semtech), M. Luis (Semtech), T. Eirich (IBM), T. Kramp (IBM),
O.Hersent (Actility)

Version: V1.0.2

Date: 2016 July

Status: Final

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17

1	Contents	
2	1 Introduction	6
3	1.1 LoRaWAN Classes	6
4	1.2 Conventions	7
5	2 Introduction on LoRaWAN options	8
6	2.1 LoRaWAN Classes	8
7	2.2 Specification scope	9
8	Class A – All end-devices	10
9	3 Physical Message Formats	11
10	3.1 Uplink Messages.....	11
11	3.2 Downlink Messages	11
12	3.3 Receive Windows.....	11
13	3.3.1 First receive window channel, data rate, and start	12
14	3.3.2 Second receive window channel, data rate, and start	12
15	3.3.3 Receive window duration.....	12
16	3.3.4 Receiver activity during the receive windows.....	12
17	3.3.5 Network sending a message to an end-device	12
18	3.3.6 Important notice on receive windows	13
19	3.3.7 Receiving or transmitting other protocols.....	13
20	4 MAC Message Formats	14
21	4.1 MAC Layer (PHYPayload).....	14
22	4.2 MAC Header (MHDR field).....	14
23	4.2.1 Message type (MType bit field).....	15
24	4.2.2 Major version of data message (Major bit field)	15
25	4.3 MAC Payload of Data Messages (MACPayload).....	16
26	4.3.1 Frame header (FHDR).....	16
27	4.3.2 Port field (FPort).....	19
28	4.3.3 MAC Frame Payload Encryption (FRMPayload).....	20
29	4.4 Message Integrity Code (MIC).....	20
30	5 MAC Commands.....	21
31	5.1 Link Check commands (<i>LinkCheckReq</i> , <i>LinkCheckAns</i>)	22
32	5.2 Link ADR commands (<i>LinkADRReq</i> , <i>LinkADRAns</i>).....	22
33	5.3 End-Device Transmit Duty Cycle (<i>DutyCycleReq</i> , <i>DutyCycleAns</i>).....	25
34	5.4 Receive Windows Parameters (<i>RXParamSetupReq</i> , <i>RXParamSetupAns</i>)	25
35	5.5 End-Device Status (<i>DevStatusReq</i> , <i>DevStatusAns</i>)	26
36	5.6 Creation / Modification of a Channel (<i>NewChannelReq</i> , <i>NewChannelAns</i> ,	
37	<i>DIChannelReq</i> , <i>DIChannelAns</i>).....	26
38	5.7 Setting delay between TX and RX (<i>RXTimingSetupReq</i> , <i>RXTimingSetupAns</i>)	29
39	5.8 End-device transmission parameters (<i>TxParamSetupReq</i> , <i>TxParamSetupAns</i>)	29
40	6 End-Device Activation	31
41	6.1 Data Stored in the End-device after Activation	31
42	6.1.1 End-device address (DevAddr).....	31
43	6.1.2 Application identifier (AppEUI).....	31
44	6.1.3 Network session key (NwkSKey)	31
45	6.1.4 Application session key (AppSKey)	31
46	6.2 Over-the-Air Activation	32
47	6.2.1 End-device identifier (DevEUI)	32
48	6.2.2 Application key (AppKey)	32
49	6.2.3 Join procedure.....	32
50	6.2.4 Join-request message	32

1	6.2.5	Join-accept message.....	33
2	6.3	Activation by Personalization	35
3	7	Retransmissions back-off.....	36
4		Class B – Beacon	37
5	8	Introduction to Class B.....	38
6	9	Principle of synchronous network initiated downlink (Class-B option).....	39
7	10	Uplink frame in Class B mode.....	41
8	11	Downlink Ping frame format (Class B option)	42
9	11.1	Physical frame format	42
10	11.2	Unicast & Multicast MAC messages.....	42
11	11.2.1	Unicast MAC message format	42
12	11.2.2	Multicast MAC message format.....	42
13	12	Beacon acquisition and tracking.....	43
14	12.1	Minimal beacon-less operation time	43
15	12.2	Extension of beacon-less operation upon reception	43
16	12.3	Minimizing timing drift.....	43
17	13	Class B Downlink slot timing	44
18	13.1	Definitions	44
19	13.2	Slot randomization	45
20	14	Class B MAC commands	46
21	14.1	PingSlotInfoReq	46
22	14.2	BeaconFreqReq.....	47
23	14.3	PingSlotChannelReq.....	47
24	14.4	BeaconTimingReq.....	48
25	14.5	BeaconTimingAns.....	49
26	15	Beaconing (Class B option).....	50
27	15.1	Beacon physical layer	50
28	15.1.1	EU 863-870MHz ISM Band	50
29	15.1.2	US 902-928MHz ISM Band	50
30	15.2	Beacon frame content	51
31	15.3	Beacon GwSpecific field format.....	52
32	15.3.1	Gateway GPS coordinate:InfoDesc = 0, 1 or 2	53
33	15.4	Beaconing precise timing	53
34	15.5	Network downlink route update requirements.....	53
35	16	Class B unicast & multicast downlink channel frequencies.....	55
36	16.1	EU 863-870MHz ISM Band	55
37	16.2	US 902-928MHz ISM Band	55
38		Class C – Continuously listening	56
39	17	Class C: Continuously listening end-device.....	57
40	17.1	Second receive window duration for Class C	57
41	17.2	Class C Multicast downlinks.....	57
42		Support information.....	59
43	18	Examples and Application Information	60
44	18.1	Uplink Timing Diagram for Confirmed Data Messages	60
45	18.2	Downlink Diagram for Confirmed Data Messages	60
46	18.3	Downlink Timing for Frame-Pending Messages	61
47	18.4	Data-Rate Adaptation during Message Retransmissions.....	62
48	19	Recommendation on contract to be provided to the network server by the end- device provider at the time of provisioning	64
49	20	Recommendation on finding the locally used channels	65
50	21	Revisions	66
51	21.1	Revision 1.0	66
52	21.2	Revision 1.0.1	66
53			

1	21.3 Revision 1.0.2	66
2	22 Glossary	67
3	23 Bibliography	68
4	23.1 References.....	68
5	24 NOTICE OF USE AND DISCLOSURE.....	69
6		

7 Tables

8	Table 1: MAC message types	15
9	Table 2: Major list.....	15
10	Table 3: FPort list.....	20
11	Table 4: MAC commands.....	22
12	Table 5: Channel state table	23
13	Table 6: LinkADRAns status bits signification	24
14	Table 7: RX2SetupAns status bits signification	26
15	Table 8: Battery level decoding	26
16	Table 9: NewChannelAns status bits signification	28
17	Table 10: DICHannelAns status bits signification	29
18	Table 11: Del mapping table	29
19	Table 11: Beacon timing	44

20

21 Figures

22	Figure 1: LoRaWAN Classes	8
23	Figure 2: Uplink PHY structure	11
24	Figure 3: Downlink PHY structure	11
25	Figure 4: End-device receive slot timing.....	12
26	Figure 5: Radio PHY structure (CRC* is only available on uplink messages)	14
27	Figure 6: PHY payload structure	14
28	Figure 7: MAC payload structure.....	14
29	Figure 8: Frame header structure.....	14
30	Figure 9: LoRa message format elements.....	14
31	Figure 10: Beacon reception slot and ping slots	40
32	Figure 11 : beacon-less temporary operation	43
33	Figure 12: Beacon timing	44
34	Figure 13: Class C end-device reception slot timing.....	57
35	Figure 14: Uplink timing diagram for confirmed data messages	60
36	Figure 15: Downlink timing diagram for confirmed data messages.....	61
37	Figure 16: Downlink timing diagram for frame-pending messages, example 1	61
38	Figure 17: Downlink timing diagram for frame-pending messages, example 2	62
39	Figure 18: Downlink timing diagram for frame-pending messages, example 3	62
40		

1 Introduction

This document describes the LoRaWAN™ network protocol which is optimized for battery-powered end-devices that may be either mobile or mounted at a fixed location.

LoRaWAN networks typically are laid out in a star-of-stars topology in which **gateways**¹ relay messages between **end-devices**² and a central **network server** at the backend. Gateways are connected to the network server via standard IP connections while end-devices use single-hop LoRa™ or FSK communication to one or many gateways.³ All communication is generally bi-directional, although uplink communication from an end-device to the network server is expected to be the predominant traffic.

Communication between end-devices and gateways is spread out on different **frequency channels** and **data rates**. The selection of the data rate is a trade-off between communication range and message duration, communications with different data rates do not interfere with each other. LoRa data rates range from 0.3 kbps to 50 kbps. To maximize both battery life of the end-devices and overall network capacity, the LoRa network infrastructure can manage the data rate and RF output for each end-device individually by means of an **adaptive data rate** (ADR) scheme.

End-devices may transmit on any channel available at any time, using any available data rate, as long as the following rules are respected:

- The end-device changes channel in a pseudo-random fashion for every transmission. The resulting frequency diversity makes the system more robust to interferences.
- The end-device respects the maximum transmit duty cycle relative to the sub-band used and local regulations.
- The end-device respects the maximum transmit duration (or dwell time) relative to the sub-band used and local regulations.

While this document specifies the protocol details, various operational parameters that are based on the regional regulations, such as maximum transmit duty-cycle and dwell time per sub-band, are described in a separate document (LoRaWAN Regional Parameters [PARAMS]). This document separation allows addition of new regional parameters without having to modify the base protocol specification.

1.1 LoRaWAN Classes

All LoRaWAN devices implement at least the Class A functionality as described in this document. In addition they may implement options named Class B, Class C as also described in this document or others to be defined. In all cases, they must remain compatible with Class A.

¹ Gateways are also known as **concentrators** or **base stations**.

² End-devices are also known as **nodes**.

³ Support for intermediate elements – repeaters – is not described in the document, however payload restrictions for encapsulation overhead are included in this specification. A repeater is defined as using LoRaWAN as its backhaul mechanism.

1 1.2 Conventions

2 MAC commands are written ***LinkCheckReq***, bits and bit fields are written **FRMPayload**,
3 constants are written RECEIVE_DELAY1, variables are written *N*.

4 In this document,

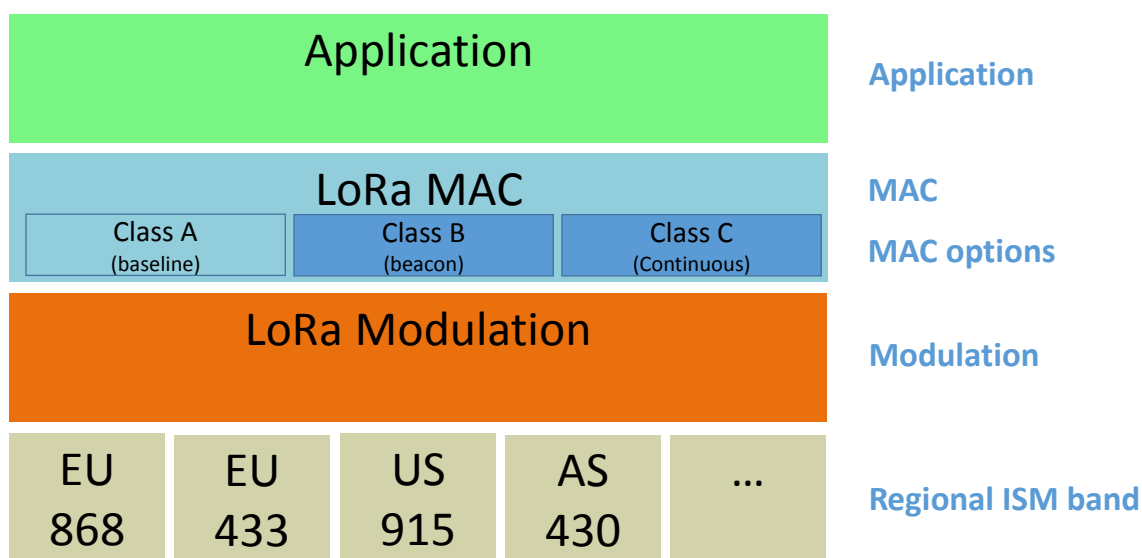
- 5 • The octet order for all multi-octet fields is little endian and
- 6 • EUI are 8 bytes multi-octet fields and are transmitted as little endian.
- 7 • By default, RFU bits are set to zero

2 Introduction on LoRaWAN options

LoRa™ is a wireless modulation for long-range low-power low-data-rate applications developed by Semtech. Devices implementing more than Class A are generally named “higher Class end-devices” in this document.

2.1 LoRaWAN Classes

A LoRa network distinguishes between a basic LoRaWAN (named Class A) and optional features (Class B, Class C ...):



8
9

Figure 1: LoRaWAN Classes

- Bi-directional end-devices (Class A):** End-devices of Class A allow for bi-directional communications whereby each end-device’s uplink transmission is followed by two short downlink receive windows. The transmission slot scheduled by the end-device is based on its own communication needs with a small variation based on a random time basis (ALOHA-type of protocol). This Class A operation is the lowest power end-device system for applications that only require downlink communication from the server shortly after the end-device has sent an uplink transmission. Downlink communications from the server at any other time will have to wait until the next scheduled uplink.
- Bi-directional end-devices with scheduled receive slots (Class B):** End-devices of Class B allow for more receive slots. In addition to the Class A random receive windows, Class B devices open extra receive windows at scheduled times. In order for the End-device to open its receive window at the scheduled time it receives a time synchronized Beacon from the gateway. This allows the server to know when the end-device is listening.
- Bi-directional end-devices with maximal receive slots (Class C):** End-devices of Class C have nearly continuously open receive windows, only closed when transmitting. Class C end-device will use more power to operate than Class A or Class B but they offer the lowest latency for server to end-device communication.

1 **2.2 Specification scope**

2 This LoRaWAN specification describes the additional functions differentiating an end-device
3 higher Class from one of Class A. A higher Class end-device shall also implement all the
4 functionality described in the LoRaWAN Class A specification.

5 **NOTE:** Physical message format, MAC message format, and other
6 parts of this specification that are common to both end-devices of
7 Class A and higher Classes are described only in the LoRaWAN
8 Class A specification to avoid redundancy.

1 **CLASS A – ALL END-DEVICES**

- 2 All LoRaWAN end-devices must implement Class A features.

1 **3 Physical Message Formats**

2 The LoRa terminology distinguishes between uplink and downlink messages.


3 **3.1 Uplink Messages**

4 **Uplink messages** are sent by end-devices to the network server relayed by one or many
5 gateways.

6 Uplink messages use the LoRa radio packet explicit mode in which the LoRa physical
7 header (**PHDR**) plus a header CRC (**PHDR_CRC**) are included.¹ The integrity of the payload
8 is protected by a CRC.

9 The **PHDR**, **PHDR_CRC** and payload **CRC** fields are inserted by the radio transceiver.

10 *Uplink PHY:*



11 **Figure 2: Uplink PHY structure**

12 **3.2 Downlink Messages**

13 Each **downlink message** is sent by the network server to only one end-device and is
14 relayed by a single gateway.²

15 Downlink messages use the radio packet explicit mode in which the LoRa physical header
16 (**PHDR**) and a header CRC (**PHDR_CRC**) are included.³

17 *Downlink PHY:*



18 **Figure 3: Downlink PHY structure**

19 **3.3 Receive Windows**

20 Following each uplink transmission the end-device opens two short receive windows. The
21 receive window start times are defined using the end of the transmission as a reference.

22

¹ See the LoRa radio transceiver datasheet for a description of LoRa radio packet implicit/explicit modes.

² This specification does not describe the transmission of multicast messages from a network server to many end-devices.

³ No payload integrity check is done at this level to keep messages as short as possible with minimum impact on any duty-cycle limitations of the ISM bands used.




Figure 4: End-device receive slot timing.

1
2

3 **3.3.1 First receive window channel, data rate, and start**

4 The first receive window RX1 uses a frequency that is a function of the uplink frequency and
 5 a data rate that is a function of the data rate used for the uplink. RX1 opens
 6 RECEIVE_DELAY1¹ seconds (+/- 20 microseconds) after the end of the uplink modulation.
 7 The relationship between uplink and RX1 slot downlink data rate is region specific and
 8 detailed in the LoRaWAN Regional Parameters document [PARAMS]. By default the first
 9 receive window datarate is identical to the datarate of the last uplink.

10 **3.3.2 Second receive window channel, data rate, and start**

11 The second receive window RX2 uses a fixed configurable frequency and data rate and
 12 opens RECEIVE_DELAY2¹ seconds (+/- 20 microseconds) after the end of the uplink
 13 modulation. The frequency and data rate used can be modified through MAC commands
 14 (see Section 5). The default frequency and data rate to use are region specific and detailed
 15 in the LoRaWAN Regional Parameters document [PARAMS].

16 **3.3.3 Receive window duration**

17 The length of a receive window must be at least the time required by the end-device's radio
 18 transceiver to effectively detect a downlink preamble.

19 **3.3.4 Receiver activity during the receive windows**

20 If a preamble is detected during one of the receive windows, the radio receiver stays active
 21 until the downlink frame is demodulated. If a frame was detected and subsequently
 22 demodulated during the first receive window and the frame was intended for this end-device
 23 after address and MIC (message integrity code) checks, the end-device does not open the
 24 second receive window.

25 **3.3.5 Network sending a message to an end-device**

26 If the network intends to transmit a downlink to an end-device, it will always initiate the
 27 transmission precisely at the beginning of one of those two receive windows.

¹ RECEIVE_DELAY1 and RECEIVE_DELAY2 are described in Chapter 6.

1 **3.3.6 Important notice on receive windows**

2 An end-device shall not transmit another uplink message before it either has received a
3 downlink message in the first or second receive window of the previous transmission, or the
4 second receive window of the previous transmission is expired.

5 **3.3.7 Receiving or transmitting other protocols**

6 The node may listen or transmit other protocols or do any transactions between the
7 LoRaWAN transmission and reception windows, as long as the end-device remains
8 compatible with the local regulation and compliant with the LoRaWAN specification.

4 MAC Message Formats


All LoRa uplink and downlink messages carry a PHY payload (**Payload**) starting with a single-octet MAC header (**MHDR**), followed by a MAC payload (**MACPayload**)¹, and ending with a 4-octet message integrity code (**MIC**).

Radio PHY layer:




Figure 5: Radio PHY structure (CRC* is only available on uplink messages)

PHYPayload:



or



or




Figure 6: PHY payload structure

MACPayload:




Figure 7: MAC payload structure

FHDR:




Figure 8: Frame header structure

Figure 9: LoRa message format elements

4.1 MAC Layer (PHYPayload)

Size (bytes)	1	1..M	4
PHYPayload	MHDR	MACPayload	MIC

The maximum length (M) of the **MACPayload** field is region specific and is specified in Chapter 6.

4.2 MAC Header (MHDR field)

Bit#	7..5	4..2	1..0
MHDR bits	MType	RFU	Major

¹ Maximum payload size is detailed in the Chapter 6.

1 The MAC header specifies the message type (**MType**) and according to which major version
 2 (**Major**) of the frame format of the LoRaWAN layer specification the frame has been
 3 encoded.

4 **4.2.1 Message type (MType bit field)**

5 The LoRaWAN distinguishes between six different MAC message types: **join request**, **join**
 6 **accept**, **unconfirmed data up/down**, and **confirmed data up/down**.

MType	Description
000	Join Request
001	Join Accept
010	Unconfirmed Data Up
011	Unconfirmed Data Down
100	Confirmed Data Up
101	Confirmed Data Down
110	RFU
111	Proprietary

7 **Table 1: MAC message types**

8 **4.2.1.1 Join-request and join-accept messages**

9 The join-request and join-accept messages are used by the over-the-air activation procedure
 10 described in Chapter 6.2.

11 **4.2.1.2 Data messages**

12 Data messages are used to transfer both MAC commands and application data, which can
 13 be combined together in a single message. A **confirmed-data message** has to be
 14 acknowledged by the receiver, whereas an **unconfirmed-data message** does not require
 15 an acknowledgment.¹ **Proprietary messages** can be used to implement non-standard
 16 message formats that are not interoperable with standard messages but must only be used
 17 among devices that have a common understanding of the proprietary extensions.

18 Message integrity is ensured in different ways for different message types and is described
 19 per message type below.

20 **4.2.2 Major version of data message (Major bit field)**

Major bits	Description
00	LoRaWAN R1
01..11	RFU

22 **Table 2: Major list**

23 **Note:** The Major version specifies the format of the messages
 24 exchanged in the join procedure (see Chapter 6.2) and the first four
 25 bytes of the MAC Payload as described in Chapter 4. For each major
 26 version, end-devices may implement different minor versions of the
 27 frame format. The minor version used by an end-device must be made
 28 known to the network server beforehand using out of band messages
 29 (e.g., as part of the device personalization information).

¹ A detailed timing diagram of the acknowledge mechanism is given in Section 18.

4.3 MAC Payload of Data Messages (MACPayload)

The MAC payload of the data messages, a so-called “data frame”, contains a frame header (FHDR) followed by an optional port field (FPort) and an optional frame payload field (FRMPayload).

4.3.1 Frame header (FHDR)

The FHDR contains the short device address of the end-device (DevAddr), a frame control octet (FCtrl), a 2-octets frame counter (FCnt), and up to 15 octets of frame options (FOpts) used to transport MAC commands.

Size (bytes)	4	1	2	0..15
FHDR	DevAddr	FCtrl	FCnt	FOpts

For downlink frames the FCtrl content of the frame header is:

Bit#	7	6	5	4	[3..0]
FCtrl bits	ADR	RFU	ACK	FPending	FOptsLen

For uplink frames the FCtrl content of the frame header is:

Bit#	7	6	5	4	[3..0]
FCtrl bits	ADR	ADRACKReq	ACK	RFU	FOptsLen

4.3.1.1 Adaptive data rate control in frame header (ADR, ADRACKReq in FCtrl)

LoRa network allows the end-devices to individually use any of the possible data rates. This feature is used by the LoRaWAN to adapt and optimize the data rate of static end-devices. This is referred to as Adaptive Data Rate (ADR) and when this is enabled the network will be optimized to use the fastest data rate possible.

Adaptive Data Rate control may not be possible when the radio channel attenuation changes fast and constantly. When the network is unable to control the data rate of a device, the device’s application layer should control it. It is recommended to use a variety of different data rates in this case. The application layer should always try to minimize the aggregated air time used given the network conditions.

If the **ADR** bit is set, the network will control the data rate of the end-device through the appropriate MAC commands. If the **ADR** bit is not set, the network will not attempt to control the data rate of the end-device regardless of the received signal quality. The **ADR** bit may be set and unset by the end-device or the Network on demand. However, whenever possible, the ADR scheme should be enabled to increase the battery life of the end-device and maximize the network capacity.

Note: Even mobile end-devices are actually immobile most of the time. So depending on its state of mobility, an end-device can request the network to optimize its data rate using ADR.

If an end-device whose data rate is optimized by the network to use a data rate higher than its lowest available data rate, it periodically needs to validate that the network still receives the uplink frames. Each time the uplink frame counter is incremented (for each new uplink, repeated transmissions do not increase the counter), the device increments an ADR_ACK_CNT counter. After ADR_ACK_LIMIT uplinks (ADR_ACK_CNT >= ADR_ACK_LIMIT) without any downlink response, it sets the ADR acknowledgment request bit (**ADRACKReq**). The network is required to respond with a downlink frame within the next ADR_ACK_DELAY frames, any received downlink frame following an uplink frame

1 resets the `ADR_ACK_CNT` counter. The downlink **ACK** bit does not need to be set as any
 2 response during the receive slot of the end-device indicates that the gateway has still
 3 received the uplinks from this device. If no reply is received within the next
 4 `ADR_ACK_DELAY` uplinks (i.e., after a total of `ADR_ACK_LIMIT + ADR_ACK_DELAY`), the
 5 end-device may try to regain connectivity by switching to the next lower data rate that
 6 provides a longer radio range. The end-device will further lower its data rate step by step
 7 every time `ADR_ACK_DELAY` is reached. The **ADRACKReq** shall not be set if the device
 8 uses its lowest available data rate because in that case no action can be taken to improve
 9 the link range.

10 **Note:** Not requesting an immediate response to an ADR
 11 acknowledgement request provides flexibility to the network to
 12 optimally schedule its downlinks.
 13

14 **Note:** In uplink transmissions the **ADRACKReq** bit is set if
 15 `ADR_ACK_CNT >= ADR_ACK_LIMIT` and the current data-rate is
 16 greater than the device defined minimum data rate, it is cleared in
 17 other conditions.

18 4.3.1.2 Message acknowledge bit and acknowledgement procedure (ACK in FCtrl)

19 When receiving a *confirmed data* message, the receiver shall respond with a data frame that
 20 has the acknowledgment bit (**ACK**) set. If the sender is an end-device, the network will send
 21 the acknowledgement using one of the receive windows opened by the end-device after the
 22 send operation. If the sender is a gateway, the end-device transmits an acknowledgment at
 23 its own discretion.

24 Acknowledgements are only sent in response to the latest message received and are never
 25 retransmitted.

26
 27 **Note:** To allow the end-devices to be as simple as possible and have
 28 as few states as possible it may transmit an explicit (possibly empty)
 29 acknowledgement data message immediately after the reception of a
 30 data message requiring a confirmation. Alternatively the end-device
 31 may defer the transmission of an acknowledgement to piggyback it
 32 with its next data message.

33 4.3.1.3 Retransmission procedure

34 The number of retransmissions (and their timing) for the same message where an
 35 acknowledgment is requested but not received is at the discretion of the end device and
 36 may be different for each end-device.

37 **Note:** Some example timing diagrams of the acknowledge mechanism
 38 are given in Chapter 18.

39
 40 **Note:** If an end-device has reached its maximum number of
 41 retransmissions without receiving an acknowledgment, it can try to re-
 42 gain connectivity by moving to a lower data rate with longer reach. It is
 43 up to the end-device to retransmit the message again or to forfeit that
 44 message and move on.

45

1 **Note:** If the network server has reached its maximum number of
 2 retransmissions without receiving an acknowledgment, it will generally
 3 consider the end-device as unreachable until it receives further
 4 messages from the end-device. It is up to the network server to
 5 retransmit the message once connectivity to the end-device in question
 6 is regained or to forfeit that message and move on.

7
 8 **Note:** The recommended data rate back-off strategy during re-
 9 transmissions is described in Chapter 18.4

10 4.3.1.4 Frame pending bit (**FPending** in **FCtrl**, downlink only)

11 The frame pending bit (**FPending**) is only used in downlink communication, indicating that
 12 the gateway has more data pending to be sent and therefore asking the end-device to open
 13 another receive window as soon as possible by sending another uplink message.

14 The exact use of **FPending** bit is described in Chapter 18.3.

15 4.3.1.5 Frame counter (**FCnt**)

16 Each end-device has two frame counters to keep track of the number of data frames sent
 17 uplink to the network server (**FCntUp**), incremented by the end-device and received by the
 18 end-device downlink from the network server (**FCntDown**), which is incremented by the
 19 network server. The network server tracks the uplink frame counter and generates the
 20 downlink counter for each end-device. After a JoinReq – JoinAccept message exchange or
 21 a reset for a personalized end-device, the frame counters on the end-device and the frame
 22 counters on the network server for that end-device are reset to 0. Subsequently **FCntUp** and
 23 **FCntDown** are incremented at the sender side by 1 for each new data frame sent in the
 24 respective direction. At the receiver side, the corresponding counter is kept in sync with the
 25 value received provided the value received has incremented compared to the current
 26 counter value and is less than the value specified by **MAX_FCNT_GAP**¹ after considering
 27 counter rollovers. If this difference is greater than the value of **MAX_FCNT_GAP** then too
 28 many data frames have been lost then subsequent will be discarded. The **FCnt** is not
 29 incremented in case of multiple transmissions of an unconfirmed frame (see **NbTrans**
 30 parameter), or in the case of a confirmed frame that is not acknowledged.

31 The LoRaWAN allows the use of either 16-bits or 32-bits frame counters. The network side
 32 needs to be informed out-of-band about the width of the frame counter implemented in a
 33 given end-device. If a 16-bits frame counter is used, the **FCnt** field can be used directly as
 34 the counter value, possibly extended by leading zero octets if required. If a 32-bits frame
 35 counter is used, the **FCnt** field corresponds to the least-significant 16 bits of the 32-bits
 36 frame counter (i.e., **FCntUp** for data frames sent uplink and **FCntDown** for data frames sent
 37 downlink).

38 The end-device shall not reuse the same **FCntUp** value, except for retransmission, with the
 39 same application and network session keys.

¹ Actual value for **MAX_FCNT_GAP**, **RECEIVE_DELAY1** and **RECEIVE_DELAY2** can be found at
Error! Reference source not found. for EU863-870 or **Error! Reference source not found.** for
 US902-928.

1 **Note:** Since the **FCnt** field carries only the least-significant 16 bits of
 2 the 32-bits frame counter, the server must infer the 16 most-significant
 3 bits of the frame counter from the observation of the traffic.

4 **4.3.1.6 Frame options (FOptsLen in FCtrl, FOpts)**

5 The frame-options length field (**FOptsLen**) in **FCtrl** byte denotes the actual length of the
 6 frame options field (**FOpts**) included in the frame.

7 **FOpts** transport MAC commands of a maximum length of 15 octets that are piggybacked
 8 onto data frames; see Chapter 5 for a list of valid MAC commands.

9 If **FOptsLen** is 0, the **FOpts** field is absent. If **FOptsLen** is different from 0, i.e. if MAC
 10 commands are present in the **FOpts** field, the port 0 cannot be used (**FPort** must be either
 11 not present or different from 0).

12 MAC commands cannot be simultaneously present in the payload field and the frame
 13 options field. Should this occur, the device shall ignore the frame.

14 **4.3.2 Port field (FPort)**

15 If the frame payload field is not empty, the port field must be present. If present, an **FPort**
 16 value of 0 indicates that the **FRMPayload** contains MAC commands only; see Chapter 4.4
 17 for a list of valid MAC commands. **FPort** values 1..223 (0x01..0xDF) are application-
 18 specific. FPort value 224 is dedicated to LoRaWAN Mac layer test protocol.

19 **Note :** The purpose of Fport value 224 is to provide a dedicated Fport
 20 to run Mac compliance test scenarios over-the-air on final versions of
 21 devices, without having to rely on specific test versions of devices for
 22 practical aspects. The test is not supposed to be simultaneous with live
 23 operations, but the Mac layer implementation of the device shall be
 24 exactly the one used for the normal application. The test protocol is
 25 normally encrypted using the AppSKey. This ensures that the network
 26 cannot enable the device's test mode without involving the device's
 27 owner. If the test runs on a live network connected device, the way the
 28 test application on the network side learns the AppSkey is outside of
 29 the scope of the LoRaWAN specification. If the test runs using OTAA
 30 on a dedicated test bench (not a live network), the way the Appkey is
 31 communicated to the test bench, for secured JOIN process, is also
 32 outside of the scope of the specification.

33 The test protocol, running at application layer, is defined outside of the
 34 LoRaWAN spec, as it is an application layer protocol.

35

36 **FPort** values 225..255 (0xE1..0xFF) are reserved for future standardized application
 37 extensions.

38

Size (bytes)	7..22	0..1	0..N
MACPayload	FHDR	FPort	FRMPayload

39 N is the number of octets of the application payload. The valid range for N is region specific
 40 and is defined in the LoRaWAN Regional Parameters document [PARAMS].

41 N should be equal or smaller than:

42
$$N \leq M - 1 - (\text{length of FHDR in octets})$$

1 where M is the maximum MAC payload length.

2 4.3.3 MAC Frame Payload Encryption (FRMPayload)

3 If a data frame carries a payload, **FRMPayload** must be encrypted before the message
4 integrity code (**MIC**) is calculated.

5 The encryption scheme used is based on the generic algorithm described in IEEE
6 802.15.4/2006 Annex B [IEEE802154] using AES with a key length of 128 bits.

7 The key K used depends on the FPort of the data message:

8

FPort	K
0	NwkSKey
1..255	AppSKey

9

Table 3: FPort list

10 The fields encrypted are:

11 $pld = \mathbf{FRMPayload}$

12 For each data message, the algorithm defines a sequence of Blocks A_i for $i = 1..k$ with $k =$
13 $\text{ceil}(\text{len}(pld) / 16)$:

Size (bytes)	1	4	1	4	4	1	1
A_i	0x01	4 x 0x00	Dir	DevAddr	FCntUp or FCntDown	0x00	i

14 The direction field (**Dir**) is 0 for uplink frames and 1 for downlink frames.

15 The blocks A_i are encrypted to get a sequence S of blocks S_i :

16

17 $S_i = \text{aes128_encrypt}(K, A_i)$ for $i = 1..k$

18 $S = S_1 | S_2 | .. | S_k$

19 Encryption and decryption of the payload is done by truncating

20

21 $(pld | \text{pad}_{16}) \text{ xor } S$

22 to the first $\text{len}(pld)$ octets.

23 4.4 Message Integrity Code (MIC)

24 The message integrity code (**MIC**) is calculated over all the fields in the message.

25

26 $msg = \mathbf{MHDR} | \mathbf{FHDR} | \mathbf{FPort} | \mathbf{FRMPayload}$

27 whereby $\text{len}(msg)$ denotes the length of the message in octets.

28 The **MIC** is calculated as follows [RFC4493]:

29

30 $cmac = \text{aes128_cmac}(\mathbf{NwkSKey}, B_0 | msg)$

31 $\mathbf{MIC} = cmac[0..3]$

32

33 whereby the block B_0 is defined as follows:

Size (bytes)	1	4	1	4	4	1	1
B_0	0x49	4 x 0x00	Dir	DevAddr	FCntUp or FCntDown	0x00	$\text{len}(msg)$

- 1
- 2 The direction field (**Dir**) is 0 for uplink frames and 1 for downlink frames.

5 MAC Commands

For network administration, a set of MAC commands may be exchanged exclusively between the network server and the MAC layer on an end-device. MAC layer commands are never visible to the application or the application server or the application running on the end-device.

A single data frame can contain any sequence of MAC commands, either piggybacked in the **FOpts** field or, when sent as a separate data frame, in the **FRMPayload** field with the **FPort** field being set to 0. Piggybacked MAC commands are always sent without encryption and must not exceed 15 octets. MAC commands sent as **FRMPayload** are always encrypted and must not exceed the maximum **FRMPayload** length.

Note: MAC commands whose content shall not be disclosed to an eavesdropper must be sent in the **FRMPayload** of a separate data message.

A MAC command consists of a command identifier (**CID**) of 1 octet followed by a possibly empty command-specific sequence of octets.

CID	Command	Transmitted by		Short Description
		End-device	Gateway	
0x02	<i>LinkCheckReq</i>	x		Used by an end-device to validate its connectivity to a network.
0x02	<i>LinkCheckAns</i>		x	Answer to LinkCheckReq command. Contains the received signal power estimation indicating to the end-device the quality of reception (link margin).
0x03	<i>LinkADRReq</i>		x	Requests the end-device to change data rate, transmit power, repetition rate or channel.
0x03	<i>LinkADRAns</i>	x		Acknowledges the LinkRateReq.
0x04	<i>DutyCycleReq</i>		x	Sets the maximum aggregated transmit duty-cycle of a device
0x04	<i>DutyCycleAns</i>	x		Acknowledges a DutyCycleReq command
0x05	<i>RXParamSetupReq</i>		x	Sets the reception slots parameters
0x05	<i>RXParamSetupAns</i>	x		Acknowledges a RXSetupReq command
0x06	<i>DevStatusReq</i>		x	Requests the status of the end-device
0x06	<i>DevStatusAns</i>	x		Returns the status of the end-device, namely its battery level and its demodulation margin
0x07	<i>NewChannelReq</i>		x	Creates or modifies the definition of a radio channel
0x07	<i>NewChannelAns</i>	x		Acknowledges a NewChannelReq command
0x08	<i>RXTimingSetupReq</i>		x	Sets the timing of the of the reception slots
0x08	<i>RXTimingSetupAns</i>	x		Acknowledges RXTimingSetupReq command
0x09	<i>TxParamSetupReq</i>		x	Used by the network server to set the maximum allowed dwell time and Max EIRP of end-device, based on local regulations
0x09	<i>TxParamSetupAns</i>	x		Acknowledges TxParamSetupReq command
0x0A	<i>DIChannelReq</i>		x	Modifies the definition of a downlink RX1 radio channel by shifting the downlink frequency from the uplink frequencies (i.e. creating an asymmetric channel)
0x0A	<i>DIChannelAns</i>	x		Acknowledges DIChannelReq command

CID	Command	Transmitted by		Short Description
		End-device	Gateway	
0x80 to 0xFF	Proprietary	x	x	Reserved for proprietary network command extensions

Table 4: MAC commands

Note: The length of a MAC command is not explicitly given and must be implicitly known by the MAC implementation. Therefore unknown MAC commands cannot be skipped and the first unknown MAC command terminates the processing of the MAC command sequence. It is therefore advisable to order MAC commands according to the version of the LoRaWAN specification which has introduced a MAC command for the first time. This way all MAC commands up to the version of the LoRaWAN specification implemented can be processed even in the presence of MAC commands specified only in a version of the LoRaWAN specification newer than that implemented.

Note: Any values adjusted by the network server (e.g., RX2, new or adjusted channels definitions) remain only valid until the next join of the end-device. Therefore after each successful join procedure the end-device uses the default parameters again and it is up to the network server to re-adjust the values again as needed.

5.1 Link Check commands (*LinkCheckReq*, *LinkCheckAns*)

With the **LinkCheckReq** command, an end-device may validate its connectivity with the network. The command has no payload.

When a **LinkCheckReq** is received by the network server via one or multiple gateways, it responds with a **LinkCheckAns** command.

Size (bytes)	1	1
LinkCheckAns Payload	Margin	GwCnt

The demodulation margin (**Margin**) is an 8-bit unsigned integer in the range of 0..254 indicating the link margin in dB of the last successfully received **LinkCheckReq** command. A value of “0” means that the frame was received at the demodulation floor (0 dB or no margin) while a value of “20”, for example, means that the frame reached the gateway 20 dB above the demodulation floor. Value “255” is reserved.

The gateway count (**GwCnt**) is the number of gateways that successfully received the last **LinkCheckReq** command.

5.2 Link ADR commands (*LinkADRReq*, *LinkADRAns*)

With the **LinkADRReq** command, the network server requests an end-device to perform a rate adaptation.

Size (bytes)	1	2	1

LinkADRReq Payload	DataRate_TXPower	ChMask	Redundancy
	Bits	[7:4]	[3:0]
	DataRate_TXPower	DataRate	TXPower

The requested data rate (**DataRate**) and TX output power (**TXPower**) are region-specific and are encoded as indicated in the LoRaWAN Regional Parameters document [PARAMS]. The TX output power indicated in the command is to be considered the maximum transmit power the device may operate at. An end-device will acknowledge as successful a command which specifies a higher transmit power than it is capable of using and should, in that case, operate at its maximum possible power. The channel mask (**ChMask**) encodes the channels usable for uplink access as follows with bit 0 corresponding to the LSB:

Bit#	Usable channels
0	Channel 1
1	Channel 2
..	..
15	Channel 16

Table 5: Channel state table

A bit in the **ChMask** field set to 1 means that the corresponding channel can be used for uplink transmissions if this channel allows the data rate currently used by the end-device. A bit set to 0 means the corresponding channels should be avoided.

Bits	7	[6:4]	[3:0]
Redundancy bits	RFU	ChMaskCntl	NbTrans

In the Redundancy bits the **NbTrans** field is the number of transmissions for each uplink message. This applies only to “unconfirmed” uplink frames. The default value is 1 corresponding to a single transmission of each frame. The valid range is [1:15]. If **NbTrans**==0 is received the end-device should use the default value. This field can be used by the network manager to control the redundancy of the node uplinks to obtain a given Quality of Service. The end-device performs frequency hopping as usual between repeated transmissions, it does wait after each repetition until the receive windows have expired. Whenever a downlink message is received during the RX1 slot window, it shall stop any further retransmission of the same uplink message. For class A devices, a reception in the RX2 slot has the same effect.

The channel mask control (**ChMaskCntl**) field controls the interpretation of the previously defined **ChMask** bit mask. It controls the block of 16 channels to which the **ChMask** applies. It can also be used to globally turn on or off all channels using specific modulation. This field usage is region specific and is defined in the LoRaWAN Regional Parameters document [PARAMS].

The network server may include multiple LinkAdrReq commands within a single downlink message. For the purpose of configuring the end-device channel mask, the end-device will process all contiguous LinkAdrReq messages, in the order present in the downlink message, as a single atomic block command. The end-device will accept or reject all Channel Mask controls in the contiguous block, and provide consistent Channel Mask ACK status indications for each command in the contiguous block in each LinkAdrAns message,

1 reflecting the acceptance or rejection of this atomic channel mask setting. The device will
 2 only process the DataRate, TXPower and NbTrans from the last message in the contiguous
 3 block, as these settings govern the end-device global state for these values. The end-
 4 device will provide consistent ACK status in each LinkAdrAns message reflecting the
 5 acceptance or rejection of these final settings.

6

7 The channel frequencies are region-specific and they are defined in Chapter 6. An end-
 8 device answers to a *LinkADRReq* with a *LinkADRAns* command.

9

Size (bytes)	1
LinkADRAns Payload	Status

10

11

Bits	[7:3]	2	1	0
Status bits	RFU	Power ACK	Data rate ACK	Channel mask ACK

12

13 The *LinkADRAns Status* bits have the following meaning:

14

	Bit = 0	Bit = 1
Channel mask ACK	The channel mask sent enables a yet undefined channel or the channel mask required all channels to be disabled. The command was discarded and the end-device state was not changed.	The channel mask sent was successfully interpreted. All currently defined channel states were set according to the mask.
Data rate ACK	The data rate requested is unknown to the end-device or is not possible given the channel mask provided (not supported by any of the enabled channels). The command was discarded and the end-device state was not changed.	The data rate was successfully set.
Power ACK	The requested power level is not implemented in the device. The command was discarded and the end-device state was not changed.	The power level was successfully set.

15

Table 6: LinkADRAns status bits signification

16 If any of those three bits equals 0, the command did not succeed and the node has kept the
 17 previous state.

5.3 End-Device Transmit Duty Cycle (*DutyCycleReq*, *DutyCycleAns*)

The ***DutyCycleReq*** command is used by the network coordinator to limit the maximum aggregated transmit duty cycle of an end-device. The aggregated transmit duty cycle corresponds to the transmit duty cycle over all sub-bands.

Size (bytes)	1	
DutyCycleReq Payload	DutyCyclePL	

Bits	7:4	3:0
DutyCyclePL	RFU	MaxDCycle

The maximum end-device transmit duty cycle allowed is:

$$\text{aggregated duty cycle} = \frac{1}{2^{\text{MaxDCycle}}}$$

The valid range for **MaxDutyCycle** is [0 : 15]. A value of 0 corresponds to “no duty cycle limitation” except the one set by the regional regulation.

An end-device answers to a ***DutyCycleReq*** with a ***DutyCycleAns*** command. The ***DutyCycleAns*** MAC reply does not contain any payload.

5.4 Receive Windows Parameters (*RXParamSetupReq*, *RXParamSetupAns*)

The ***RXParamSetupReq*** command allows a change to the frequency and the data rate set for the second receive window (RX2) following each uplink. The command also allows to program an offset between the uplink and the RX1 slot downlink data rates.

Size (bytes)	1	3
RXParamSetupReq Payload	DLsettings	Frequency

Bits	7	6:4	3:0
DLsettings	RFU	RX1DRoffset	RX2DataRate

The **RX1DRoffset** field sets the offset between the uplink data rate and the downlink data rate used to communicate with the end-device on the first reception slot (RX1). As a default this offset is 0. The offset is used to take into account maximum power density constraints for base stations in some regions and to balance the uplink and downlink radio link margins.

The data rate (**RX2DataRate**) field defines the data rate of a downlink using the second receive window following the same convention as the ***LinkADRReq*** command (0 means DR0/125kHz for example). The frequency (**Frequency**) field corresponds to the frequency of the channel used for the second receive window, whereby the frequency is coded following the convention defined in the ***NewChannelReq*** command.

The ***RXParamSetupAns*** command is used by the end-device to acknowledge the reception of ***RXParamSetupReq*** command. The ***RXParamSetupAns*** command should be added in the FOpt field of all uplinks until a class A downlink is received by the end-device. This guarantees that even in presence of uplink packet loss, the network is always aware of the downlink parameters used by the end-device.

The payload contains a single status byte.

Size (bytes)	1
RXParamSetupAns Payload	Status

1
2 The status (**Status**) bits have the following meaning.

Bits	7:3	2	1	0
Status bits	RFU	RX1DRoffset ACK	RX2 Data rate ACK	Channel ACK

	Bit = 0	Bit = 1
Channel ACK	The frequency requested is not usable by the end-device.	RX2 slot channel was successfully set
RX2 Data rate ACK	The data rate requested is unknown to the end-device.	RX2 slot data rate was successfully set
RX1DRoffset ACK	the uplink/downlink data rate offset for RX1 slot is not in the allowed range	RX1DRoffset was successfully set

3
4 **Table 7: RX2SetupAns status bits signification**

5 If either of the 3 bits is equal to 0, the command did not succeed and the previous
6 parameters are kept.

8 **5.5 End-Device Status (*DevStatusReq*, *DevStatusAns*)**

9 With the ***DevStatusReq*** command a network server may request status information from an
10 end-device. The command has no payload. If a ***DevStatusReq*** is received by an end-
11 device, it responds with a ***DevStatusAns*** command.

Size (bytes)	1	1
DevStatusAns Payload	Battery	Margin

12 The battery level (**Battery**) reported is encoded as follows:

Battery	Description
0	The end-device is connected to an external power source.
1..254	The battery level, 1 being at minimum and 254 being at maximum
255	The end-device was not able to measure the battery level.

13 **Table 8: Battery level decoding**

14 The margin (**Margin**) is the demodulation signal-to-noise ratio in dB rounded to the nearest
15 integer value for the last successfully received ***DevStatusReq*** command. It is a signed
16 integer of 6 bits with a minimum value of -32 and a maximum value of 31.

Bits	7:6	5:0
Status	RFU	Margin

17 **5.6 Creation / Modification of a Channel (*NewChannelReq*, 18 *NewChannelAns*, *DIChannelReq*, *DIChannelAns*)**

19

1 The **NewChannelReq** command can be used to either modify the parameters of an existing
 2 bidirectional channel or to create a new one. The command sets the center frequency of the
 3 new channel and the range of uplink data rates usable on this channel:

Size (bytes)	1	3	1
NewChannelReq Payload	ChIndex	Freq	DrRange

4
 5 The channel index (**ChIndex**) is the index of the channel being created or modified.
 6 Depending on the region and frequency band used, the LoRaWAN specification imposes
 7 default channels which must be common to all devices and cannot be modified by the
 8 **NewChannelReq** command (cf. Chapter 6). If the number of default channels is N , the
 9 default channels go from 0 to $N-1$, and the acceptable range for **ChIndex** is N to 15. A
 10 device must be able to handle at least 16 different channel definitions. In certain region the
 11 device may have to store more than 16 channel definitions.
 12

13 The frequency (**Freq**) field is a 24 bits unsigned integer. The actual channel frequency in Hz
 14 is $100 \times \mathbf{Freq}$ whereby values representing frequencies below 100 MHz are reserved for
 15 future use. This allows setting the frequency of a channel anywhere between 100 MHz to
 16 1.67 GHz in 100 Hz steps. A **Freq** value of 0 disables the channel. The end-device has to
 17 check that the frequency is actually allowed by its radio hardware and return an error
 18 otherwise.
 19

20 The data-rate range (**DrRange**) field specifies the uplink data-rate range allowed for this
 21 channel. The field is split in two 4-bit indexes:

Bits	7:4	3:0
DrRange	MaxDR	MinDR

22
 23 Following the convention defined in Section 5.2 the minimum data rate (**MinDR**) subfield
 24 designate the lowest uplink data rate allowed on this channel. For example 0 designates
 25 DR0 / 125 kHz. Similarly, the maximum data rate (**MaxDR**) designates the highest uplink
 26 data rate. For example, $\mathbf{DrRange} = 0x77$ means that only 50 kbps GFSK is allowed on a
 27 channel and $\mathbf{DrRange} = 0x50$ means that DR0 / 125 kHz to DR5 / 125 kHz are supported.

28 The newly defined or modified channel is enabled and can immediately be used for
 29 communication. The RX1 downlink frequency is set equal to the uplink frequency.

30 The end-device acknowledges the reception of a **NewChannelReq** by sending back a
 31 **NewChannelAns** command. The payload of this message contains the following
 32 information:

Size (bytes)	1
NewChannelAns Payload	Status

33
 34 The status (**Status**) bits have the following meaning:

Bits	7:2	1	0
Status	RFU	Data rate range ok	Channel frequency ok

	Bit = 0	Bit = 1
Data rate range ok	The designated data rate range exceeds the ones currently defined for this end-device	The data rate range is compatible with the possibilities of the end-device

Channel frequency ok	The device cannot use this frequency	The device is able to use this frequency.
-----------------------------	--------------------------------------	---

 Table 9: *NewChannelAns* status bits signification

If either of those 2 bits equals 0, the command did not succeed and the new channel has not been created.

The ***DIChannelReq*** command allows the network to associate a different downlink frequency to the RX1 slot. This command is applicable for all the geographic regions supporting the ***NewChannelReq*** command (for example EU and China, but not for US or Australia as described in the LoRaWAN Regional Parameters document [PARAMS]).

The command sets the center frequency used for the downlink RX1 slot, as follows:

Size (bytes)	1	3
<i>DIChannelReq</i> Payload	ChIndex	Freq

The channel index (**ChIndex**) is the index of the channel whose downlink frequency is modified

The frequency (**Freq**) field is a 24 bits unsigned integer. The actual downlink frequency in Hz is 100 x **Freq** whereby values representing frequencies below 100 MHz are reserved for future use. The end-device has to check that the frequency is actually allowed by its radio hardware and return an error otherwise.

The end-device acknowledges the reception of a ***DIChannelReq*** by sending back a ***DIChannelAns*** command. The ***DIChannelAns*** command shall be added in the FOpt field of all uplinks until a downlink packet is received by the end-device. This guarantees that even in presence of uplink packet loss, the network is always aware of the downlink frequencies used by the end-device.

The payload of this message contains the following information:

Size (bytes)	1
<i>DIChannelAns</i> Payload	Status

The status (**Status**) bits have the following meaning:

Bits	7:2	1	0
Status	RFU	Uplink frequency exists	Channel frequency ok

	Bit = 0	Bit = 1
Channel frequency ok	The device cannot use this frequency	The device is able to use this frequency.
Uplink frequency exists	The uplink frequency is not defined for this channel , the downlink frequency can only be set for a channel that already has a valid uplink frequency	The uplink frequency of the channel is valid

Table 10: DIChannelAns status bits signification

5.7 Setting delay between TX and RX (*RXTimingSetupReq*, *RXTimingSetupAns*)

The *RXTimingSetupReq* command allows configuring the delay between the end of the TX uplink and the opening of the first reception slot. The second reception slot opens one second after the first reception slot.

Size (bytes)	1
<i>RXTimingSetupReq</i> Payload	Settings

The delay (**Delay**) field specifies the delay. The field is split in two 4-bit indexes:

Bits	7:4	3:0
Settings	RFU	Del

The delay is expressed in seconds. **Del** 0 is mapped on 1 s.

Del	Delay [s]
0	1
1	1
2	2
3	3
..	..
15	15

Table 11: Del mapping table

An end device answers *RXTimingSetupReq* with *RXTimingSetupAns* with no payload.

The *RXTimingSetupAns* command should be added in the FOpt field of all uplinks until a class A downlink is received by the end-device. This guarantees that even in presence of uplink packet loss, the network is always aware of the downlink parameters used by the end-device.

5.8 End-device transmission parameters (*TxParamSetupReq*, *TxParamSetupAns*)

This MAC command only needs to be implemented for compliance in certain regulatory regions. Please refer to the LoRaWAN Regional Parameters [PARAMS] document.

The *TxParamSetupReq* command can be used to notify the end-device of the maximum allowed dwell time, i.e. the maximum continuous transmission time of a packet over the air, as well as the maximum allowed end-device Effective Isotropic Radiated Power (EIRP).

Size (bytes)	1
<i>TxParamSetup</i> payload	EIRP_DwellTime

The structure of EIRP_DwellTime field is described below:

Bits	7:6	5	4	3:0
MaxDwellTime	RFU	DownlinkDwellTime	UplinkDwellTime	MaxEIRP

1

2 Bits [0...3] of ***TxParamSetupReq*** command are used to encode the Max EIRP value, as per
 3 the following table. The EIRP values in this table are chosen in a way that covers a wide
 4 range of max EIRP limits imposed by the different regional regulations.

5

Coded Value	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Max EIRP (dBm)	8	10	12	13	14	16	18	20	21	24	26	27	29	30	33	36

6 The maximum EIRP corresponds to an upper bound on the device’s radio transmit power.
 7 The device is not required to transmit at that power , but shall never radiate more that this
 8 specified EIRP.

9 Bits 4 and 5 define the maximum Uplink and downlink dwell time respectively, which is
 10 encoded as per the following table:

Coded Value	Dwell Time
0	No Limit
1	400 ms

11

12 When this MAC command is implemented (region specific), the end-device acknowledges
 13 the ***TxParamSetupReq*** command by sending a ***TxParamSetupAns*** command. This
 14 ***TxParamSetupAns*** command doesn’t contain any payload.

15 When this MAC command is used in a region where it is not required, the device does not
 16 process it and shall not transmit an acknowledgement.

17

18

19

1 **6 End-Device Activation**

2 To participate in a LoRaWAN network, each end-device has to be personalized and
3 activated.

4 Activation of an end-device can be achieved in two ways, either via **Over-The-Air**
5 **Activation** (OTAA) when an end-device is deployed or reset, or via **Activation By**
6 **Personalization** (ABP) in which the two steps of end-device personalization and activation
7 are done as one step.

8 **6.1 Data Stored in the End-device after Activation**

9 After activation, the following information is stored in the end-device: a device address
10 (**DevAddr**), an application identifier (**AppEUI**), a network session key (**NwkSKey**), and an
11 application session key (**AppSKey**).

12 **6.1.1 End-device address (DevAddr)**

13 The **DevAddr** consists of 32 bits identifies the end-device within the current network. Its
14 format is as follows:

Bit#	[31..25]	[24..0]
DevAddr bits	NwkID	NwkAddr

15 The most significant 7 bits are used as network identifier (**NwkID**) to separate addresses of
16 territorially overlapping networks of different network operators and to remedy roaming
17 issues. The least significant 25 bits, the network address (**NwkAddr**) of the end-device, can
18 be arbitrarily assigned by the network manager.

19 **6.1.2 Application identifier (AppEUI)**

20 The **AppEUI** is a global application ID in IEEE EUI64 address space that uniquely identifies
21 the entity able to process the JoinReq frame.

22 The **AppEUI** is stored in the end-device before the activation procedure is executed.

23 **6.1.3 Network session key (NwkSKey)**

24 The **NwkSKey** is a **network session key** specific for the end-device. It is used by both the
25 network server and the end-device to calculate and verify the **MIC** (message integrity code)
26 of all data messages to ensure data integrity. It is further used to encrypt and decrypt the
27 payload field of a MAC-only data messages.

28 **6.1.4 Application session key (AppSKey)**

29 The **AppSKey** is an **application session key** specific for the end-device. It is used by both
30 the application server and the end-device to encrypt and decrypt the payload field of
31 application-specific data messages. Application payloads are end-to-end encrypted between
32 the end-device and the application server, but they are not integrity protected. That means, a
33 network server may be able to alter the content of the data messages in transit. Network
34 servers are considered as trusted, but applications wishing to implement end-to-end
35 confidentiality and integrity protection are recommended to use additional end-to-end
36 security solutions, which are beyond the scope of this specification.

37

1 **6.2 Over-the-Air Activation**

2 For over-the-air activation, end-devices must follow a join procedure prior to participating in
 3 data exchanges with the network server. An end-device has to go through a new join
 4 procedure every time it has lost the session context information.

5 The join procedure requires the end-device to be personalized with the following information
 6 before its starts the join procedure: a globally unique end-device identifier (**DevEUI**), the
 7 application identifier (**AppEUI**), and an AES-128 key (**AppKey**).

8 The **AppEUI** is described above in 6.1.2.

9 **Note:** For over-the-air-activation, end-devices are not personalized
 10 with any kind of network key. Instead, whenever an end-device joins a
 11 network, a network session key specific for that end-device is derived
 12 to encrypt and verify transmissions at the network level. This way,
 13 roaming of end-devices between networks of different providers is
 14 facilitated. Using both a network session key and an application
 15 session key further allows federated network servers in which
 16 application data cannot be read or tampered with by the network
 17 provider.

18 **6.2.1 End-device identifier (DevEUI)**

19 The **DevEUI** is a global end-device ID in IEEE EUI64 address space that uniquely identifies
 20 the end-device.

21 **6.2.2 Application key (AppKey)**

22 The **AppKey** is an AES-128 root key specific to the end-device.¹ Whenever an end-device
 23 joins a network via over-the-air activation, the AppKey is used to derive the session keys
 24 NwkSKey and AppSKey specific for that end-device to encrypt and verify network
 25 communication and application data.

26 **6.2.3 Join procedure**

27 From an end-device's point of view, the join procedure consists of two MAC messages
 28 exchanged with the server, namely a **join request** and a **join accept**.

29 **6.2.4 Join-request message**

30 The join procedure is always initiated from the end-device by sending a join-request
 31 message.

32

Size (bytes)	8	8	2
Join Request	AppEUI	DevEUI	DevNonce

33 The join-request message contains the **AppEUI** and **DevEUI** of the end-device followed by a
 34 **nonce** of 2 octets (**DevNonce**).

1. Since all end-devices end up with unrelated application keys specific for each end-device, extracting the AppKey from an end-device only compromises this one end-device.

1 **DevNonce** is a random value.¹ For each end-device, the network server keeps track of a
 2 certain number of **DevNonce** values used by the end-device in the past, and ignores join
 3 requests with any of these **DevNonce** values from that end-device.

4 **Note:** This mechanism prevents replay attacks by sending previously
 5 recorded join-request messages with the intention of disconnecting the
 6 respective end-device from the network. Any time the network server
 7 processes a Join-Request and generates a Join-accept frame, it shall
 8 maintain both the old security context (keys and counters, if any) and
 9 the new one until it receives the first successful uplink frame using the
 10 new context, after which the old context can be safely removed. This
 11 provides defense against an adversary replaying an earlier Join-
 12 request using a DevNonce that falls outside the finite list of values
 13 tracked by the network server.

14 The message integrity code (**MIC**) value (see Chapter 4 for MAC message description) for a
 15 join-request message is calculated as follows:²

$$17 \quad cmac = aes128_cmac(AppKey, MHDR | AppEUI | DevEUI | DevNonce)$$

$$18 \quad MIC = cmac[0..3]$$

19 The join-request message is not encrypted.

20 The join-request message can be transmitted using any data rate and following a random
 21 frequency hopping sequence across the specified join channels. It is recommended to use a
 22 plurality of data rates. The intervals between transmissions of **Join-Requests** shall respect
 23 the condition described in chapter 7.

24 6.2.5 Join-accept message

25 The network server will respond to the **join-request** message with a **join-accept** message if
 26 the end-device is permitted to join a network. The join-accept message is sent like a normal
 27 downlink but uses delays JOIN_ACCEPT_DELAY1 or JOIN_ACCEPT_DELAY2 (instead of
 28 RECEIVE_DELAY1 and RECEIVE_DELAY2, respectively). The channel frequency and data
 29 rate used for these two receive windows are identical to the one used for the RX1 and RX2
 30 receive windows described in the “receive windows” section of the LoRaWAN Regional
 31 Parameters document [PARAMS].

32 No response is given to the end-device if the join request is not accepted.

33 The join-accept message contains an application nonce (**AppNonce**) of 3 octets, a network
 34 identifier (**NetID**), an end-device address (**DevAddr**), a delay between TX and RX
 35 (**RxDelay**) and an optional list of channel frequencies (**CFList**) for the network the end-
 36 device is joining. The CFList option is region specific and is defined in the LoRaWAN
 37 Regional Parameters document [PARAMS].

38

Size (bytes)	3	3	4	1	1	(16) Optional
Join Accept	AppNonce	NetID	DevAddr	DLSettings	RxDelay	CFList

¹ The DevNonce can be extracted by issuing a sequence of RSSI measurements under the assumption that the quality of randomness fulfills the criteria of true randomness

² [RFC4493]

1 The **AppNonce** is a random value or some form of unique ID provided by the network server
 2 and used by the end-device to derive the two session keys **NwkSKey** and **AppSKey** as
 3 follows:¹

4 $NwkSKey = \text{aes128_encrypt}(\text{AppKey}, 0x01 \mid \text{AppNonce} \mid \text{NetID} \mid \text{DevNonce} \mid \text{pad}_{16})$

5 $\text{AppSKey} = \text{aes128_encrypt}(\text{AppKey}, 0x02 \mid \text{AppNonce} \mid \text{NetID} \mid \text{DevNonce} \mid \text{pad}_{16})$

6 The MIC value for a join-accept message is calculated as follows:²

7 $\text{cmac} = \text{aes128_cmac}(\text{AppKey},$

8 MHDR \mid AppNonce \mid NetID \mid DevAddr \mid DLSettings \mid RxDelay \mid CFList)

9 $\text{MIC} = \text{cmac}[0..3]$

10 The join-accept message itself is encrypted with the **AppKey** as follows:

11 $\text{aes128_decrypt}(\text{AppKey}, \text{AppNonce} \mid \text{NetID} \mid \text{DevAddr} \mid \text{DLSettings} \mid \text{RxDelay} \mid \text{CFList} \mid$

12 MIC)

13 **Note:** The network server uses an AES decrypt operation in ECB
 14 mode to encrypt the join-accept message so that the end-device can
 15 use an AES encrypt operation to decrypt the message. This way an
 16 end-device only has to implement AES encrypt but not AES decrypt.

18 **Note:** Establishing these two session keys allows for a federated
 19 network server infrastructure in which network operators are not able
 20 to eavesdrop on application data. In such a setting, the application
 21 provider must support the network operator in the process of an end-
 22 device actually joining the network and establishing the NwkSKey for
 23 the end-device. At the same time the application provider commits to
 24 the network operator that it will take the charges for any traffic incurred
 25 by the end-device and retains full control over the AppSKey used for
 26 protecting its application data.

27 The format of the **NetID** is as follows: The seven LSB of the **NetID** are called **NwkID** and
 28 match the seven MSB of the short address of an end-device as described before.
 29 Neighboring or overlapping networks must have different **NwkIDs**. The remaining 17 MSB
 30 can be freely chosen by the network operator.

31 The DLsettings field contains the downlink configuration:

Bits	7	6:4	3:0
DLsettings	RFU	RX1DRoffset	RX2 Data rate

33
 34 The RX1DRoffset field sets the offset between the uplink data rate and the downlink data
 35 rate used to communicate with the end-device on the first reception slot (RX1). By default
 36 this offset is 0.. The offset is used to take into account maximum power density constraints
 37 for base stations in some regions and to balance the uplink and downlink radio link margins.

38 The actual relationship between the uplink and downlink data rate is region specific and
 39 detailed in the LoRaWAN Regional Parameters document [PARAMS].

40 The delay **RxDelay** follows the same convention as the **Delay** field in the
 41 **RXTimingSetupReq** command.

¹ The pad₁₆ function appends zero octets so that the length of the data is a multiple of 16.

² [RFC4493]

1 **6.3 Activation by Personalization**

2 Under certain circumstances, end-devices can be activated by personalization. Activation by
3 personalization directly ties an end-device to a specific network by-passing the **join request**
4 - **join accept** procedure.

5 Activating an end-device by personalization means that the **DevAddr** and the two session
6 keys **NwkSKey** and **AppSKey** are directly stored into the end-device instead of the **DevEUI**,
7 **AppEUI** and the **AppKey**. The end-device is equipped with the required information for
8 participating in a specific LoRa network when started.

9 Each device should have a unique set of NwkSKey and AppSKey. Compromising the keys
10 of one device shouldn't compromise the security of the communications of other devices.
11 The process to build those keys should be such that the keys cannot be derived in any way
12 from publicly available information (like the node address for example).

7 Retransmissions back-off

Uplink frames that:

- Require an **acknowledgement or an answer** by the network or an application server, and are **retransmitted** by the device if the acknowledgement or answer is not received.

and

- can be triggered by an **external** event causing **synchronization** across a large (>100) number of devices (power outage, radio jamming, network outage, earthquake...)

can trigger a catastrophic, self-persisting, radio network overload situation.

Note: An example of such uplink frame is typically the JoinRequest if the implementation of a group of end-devices decides to reset the MAC layer in the case of a network outage.

The whole group of end-device will start broadcasting JoinRequest uplinks and will only stops when receiving a JoinResponse from the network.

For those frame retransmissions, the interval between the end of the RX2 slot and the next uplink retransmission shall be random and follow a different sequence for every device (For example using a pseudo-random generator seeded with the device's address) .The transmission duty-cycle of such message shall respect the local regulation and the following limits, whichever is more constraining:

Aggregated during the first hour following power-up or reset	$T_0 < t < T_0 + 1$	Transmit time < 36Sec
Aggregated during the next 10 hours	$T_0 + 1 < t < T_0 + 11$	Transmit time < 36Sec
After the first 11 hours , aggregated over 24h	$T_0 + 11 + N < t < T_0 + 35 + N$ $N \geq 0$	Transmit time < 8.7Sec per 24h

CLASS B – BEACON

1
2
3

Class B must be considered as experimental in this version of the specification

1 **8 Introduction to Class B**

2 This section describes the LoRaWAN Class B layer which is optimized for battery-powered
3 end-devices that may be either mobile or mounted at a fixed location.

4 End-devices should implement Class B operation when there is a requirement to open
5 receive windows at fixed time intervals for the purpose of enabling server initiated downlink
6 messages.

7 LoRaWAN Class B option adds a synchronized reception window on the end-device.

8 One of the limitations of LoRaWAN Class A is the Aloha method of sending data from the
9 end-device; it does not allow for a known reaction time when the customer application or the
10 server wants to address the end-device. The purpose of Class B is to have an end-device
11 available for reception on a predictable time, in addition to the reception windows that
12 follows the random uplink transmission from the end-device of Class A. Class B is achieved
13 by having the gateway sending a beacon on a regular basis to synchronize the all the end-
14 devices in the network so that the end-device can opening a short extra reception window
15 (called “ping slot”) at a predictable time during a periodic time slot.

16 **Note:** The decision to switch from Class A to Class B comes from the
17 application layer of the end-device. If this class A to Class B switch
18 needs to be controlled from the network side, the customer application
19 must use one of the end-device’s Class A uplinks to send back a
20 downlink to the application layer, and it needs the application layer on
21 the end-device to recognize this request – this process is not managed
22 at the LoRaWAN level.


9 Principle of synchronous network initiated downlink (Class-B option)

For a network to support end-devices of Class B, all gateways must synchronously broadcast a beacon providing a timing reference to the end-devices. Based on this timing reference the end-devices can periodically open receive windows, hereafter called “ping slots”, which can be used by the network infrastructure to initiate a downlink communication. A network initiated downlink using one of these ping slots is called a “ping”. The gateway chosen to initiate this downlink communication is selected by the network server based on the signal quality indicators of the last uplink of the end-device. For this reason, if an end-device moves and detects a change in the identity advertised in the received beacon, it must send an uplink to the network server so that the server can update the downlink routing path database.

All end-devices start and join the network as end-devices of Class A. The end-device application can then decide to switch to Class B. This is done through the following process:

- The end-device application requests the LoRaWAN layer to switch to Class B mode. The LoRaWAN layer in the end-device searches for a beacon and returns either a BEACON_LOCKED service primitive to the application if a network beacon was found and locked or a BEACON_NOT_FOUND service primitive. To accelerate the beacon discovery the LoRaWAN layer may use the “BeaconTimingReq” message described later.
- Based on the beacon strength and the battery life constraints, the end-device application selects a ping slot data rate and periodicity, this is then requested them from the end-device LoRaWAN layer.
- Once in Class B mode, the MAC layer sets to 1 the *Class B* bit of the FCTRL field of every uplink frame transmitted. This bit signals to the server that the device has switched to Class B. The MAC layer will autonomously schedule a reception slot for each beacon and each ping slot. When the beacon reception is successful the end-device LoRaWAN layer forwards the beacon content to the application together with the measured radio signal strength. The end-device LoRaWAN layer takes into account the maximum possible clock drift in the scheduling of the beacon reception slot and ping slots. When a downlink is successfully demodulated during a ping slot, it is processed similarly to a downlink as described in the LoRaWAN Class A specification.
- A mobile end-device must periodically inform the network server of its location to update the downlink route. This is done by transmitting a normal (possibly empty) “unconfirmed” or “confirmed” uplink. The end-device LoRaWAN layer will appropriately set the *Class B* bit to 1. Optimally this can be done more efficiently if the application detects that the node is moving by analyzing the beacon content. In that case the end-device must apply a random delay (as defined in Section 15.5 between the beacon reception and the uplink transmission to avoid systematic uplink collisions.
- If no beacon has been received for a given period (as defined in Section 12.2), the synchronization with the network is lost. The MAC layer must inform the application layer that it has switched back to Class A. As a consequence the end-device LoRaWAN layer stops setting the *Class B* bit in all uplinks and this informs the network server that the end-device is no longer in Class B mode. The end-device application can try to switch back to Class B periodically. This will restart this process starting with a beacon search.

1 The following diagram illustrates the concept of beacon reception slots and ping slots.



2
3

Figure 10: Beacon reception slot and ping slots

4 In this example, given the beacon period is 128 s, the end-device also opens a ping
 5 reception slot every 32 s. Most of the time this ping slot is not used by the server and
 6 therefore the end-device reception window is closed as soon as the radio transceiver has
 7 assessed that no preamble is present on the radio channel. If a preamble is detected the
 8 radio transceiver will stay on until the downlink frame is demodulated. The MAC layer will
 9 then process the frame, check that its address field matches the end-device address and
 10 that the Message Integrity Check is valid before forwarding it to the application layer.

1 **10 Uplink frame in Class B mode**

2 The uplink frames in Class B mode are same as the Class A uplinks with the exception of
 3 the RFU bit in the FCtrl field in the Frame header. In the Class A uplink this bit is unused
 4 (RFU). This bit is used for Class B uplinks.

5

Bit#	7	6	5	4	3..0
FCtrl	ADR	ADRACKReq	ACK	Class B	FOptsLen

6
 7 The *Class B* bit set to 1 in an uplink signals the network server that the device as switched to
 8 Class B mode and is now ready to receive scheduled downlink pings.

9
 10 The signification of the FPending bit for downlink is unaltered and still signals that one or
 11 more downlink frames are queued for this device in the server and that the device should
 12 keep is receiver on as described in the Class A specification.

13

1 **11 Downlink Ping frame format (Class B option)**

2 **11.1 Physical frame format**

3 A downlink Ping uses the same format as a Class A downlink frame but might follow a
4 different channel frequency plan.

5 **11.2 Unicast & Multicast MAC messages**

6 Messages can be “unicast” or “multicast”. Unicast messages are sent to a single end-device
7 and multicast messages are sent to multiple end-devices. All devices of a multicast group
8 must share the same multicast address and associated encryption keys. The LoRaWAN
9 Class B specification does not specify means to remotely setup such a multicast group or
10 securely distribute the required multicast key material. This must either be performed during
11 the node personalization or through the application layer.

12 **11.2.1 Unicast MAC message format**

13 The MAC payload of a unicast downlink **Ping** uses the format defined in the Class A
14 specification. It is processed by the end-device in exactly the same way. The same frame
15 counter is used and incremented whether the downlink uses a Class B ping slot or a Class A
16 “piggy-back” slot.

17 **11.2.2 Multicast MAC message format**

18 The Multicast frames share most of the unicast frame format with a few exceptions:

- 19 • They are not allowed to carry MAC commands, neither in the **FOpt** field, nor in the
20 payload on port 0 because a multicast downlink does not have the same
21 authentication robustness as a unicast frame.
- 22 • The **ACK** and **ADRACKReq** bits must be zero. The **MType** field must carry the value
23 for Unconfirmed Data Down.
- 24 • The **FPending** bit indicates there is more multicast data to be sent. If it is set the
25 next multicast receive slot will carry a data frame. If it is not set the next slot may or
26 may not carry data. This bit can be used by end-devices to evaluate priorities for
27 conflicting reception slots.

28

1 **12 Beacon acquisition and tracking**

2 Before switching from Class A to Class B, the end-device must first receive one of the
3 network beacons to align his internal timing reference with the network.

4 Once in Class B, the end-device must periodically search and receive a network beacon to
5 cancel any drift of its internal clock time base, relative to the network timing.

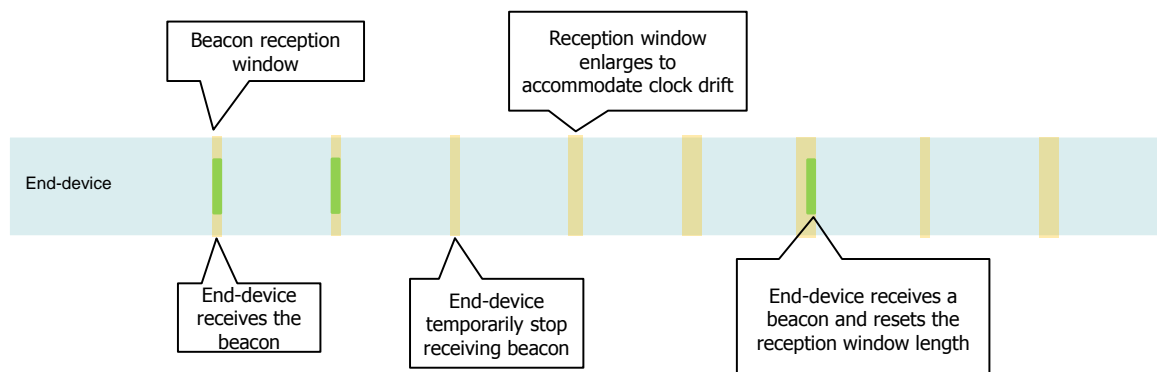
6 A Class B device may be temporarily unable to receive beacons (out of range from the
7 network gateways, presence of interference, ..). In this event, the end-device has to
8 gradually widen its beacon and ping slots reception windows to take into account a possible
9 drift of its internal clock.

10 **Note:** For example, a device which internal clock is defined with a +/-
11 10ppm precision may drift by +/-1.3mSec every beacon period.

12 **12.1 Minimal beacon-less operation time**

13 In the event of beacon loss, a device shall be capable of maintaining Class B operation for 2
14 hours (120 minutes) after it received the last beacon. This temporary Class B operation
15 without beacon is called “beacon-less” operation. It relies on the end-device’s own clock to
16 keep timing.

17 During beacon-less operation, unicast, multicast and beacon reception slots must all be
18 progressively expanded to accommodate the end-device’s possible clock drift.
19



20 **Figure 11 : beacon-less temporary operation**

22 **12.2 Extension of beacon-less operation upon reception**

23 During this 120 minutes time interval the reception of any beacon directed to the end-device,
24 should extend the Class B beacon-less operation further by another 120 minutes as it allows
25 to correct any timing drift and reset the receive slots duration.

26 **12.3 Minimizing timing drift**

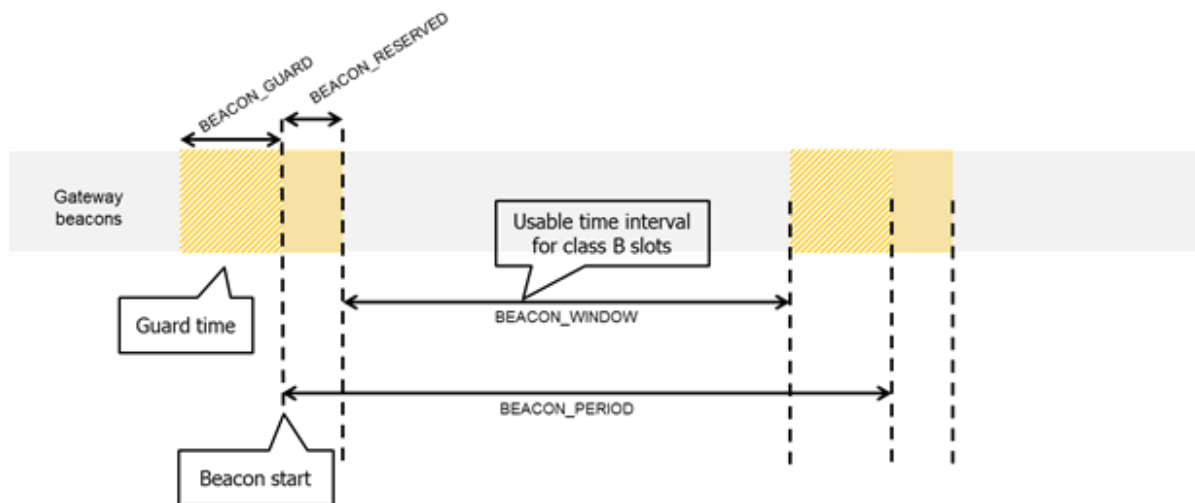
27 The end-devices may use the beacon’s (when available) precise periodicity to calibrate their
28 internal clock and therefore reduce the initial clock frequency imprecision. As the timing
29 oscillator’s exhibit a predictable temperature frequency shift, the use of a temperature
30 sensor could enable further minimization of the timing drift.

1 **13 Class B Downlink slot timing**

2 **13.1 Definitions**

3 To operate successfully in Class B the end-device must open reception slots at precise
4 instants relative to the infrastructure beacon. This section defines the required timing.

5 The interval between the start of two successive beacons is called the beacon period. The
6 beacon frame transmission is aligned with the beginning of the BEACON_RESERVED
7 interval. Each beacon is preceded by a guard time interval where no ping slot can be placed.
8 The length of the guard interval corresponds to the time on air of the longest allowed frame.
9 This is to insure that a downlink initiated during a ping slot just before the guard time will
10 always have time to complete without colliding with the beacon transmission. The usable
11 time interval for ping slot therefore spans from the end of the beacon reserved time interval
12 to the beginning of the next beacon guard interval.



13 **Figure 12: Beacon timing**

14

Beacon_period	128 s
Beacon_reserved	2.120 s
Beacon_guard	3.000 s
Beacon-window	122.880 s

15 **Table 12: Beacon timing**

16 The beacon frame time on air is actually much shorter than the beacon reserved time
17 interval to allow appending network management broadcast frames in the future.

18 The beacon window interval is divided into $2^{12} = 4096$ ping slots of 30 ms each numbered
19 from 0 to 4095.

20 An end-device using the slot number N must turn on its receiver exactly T_{on} seconds after
21 the start of the beacon where:

22
$$T_{on} = beacon_reserved + N * 30 \text{ ms}$$

23 N is called the *slot index*.

24 The latest ping slot starts at $beacon_reserved + 4095 * 30 \text{ ms} = 124\,970 \text{ ms}$ after the
25 beacon start or 3030 ms before the beginning of the next beacon.

1 13.2 Slot randomization

2 To avoid systematic collisions or over-hearing problems the slot index is randomized and
3 changed at every beacon period.

4 The following parameters are used:

5

DevAddr	Device 32 bit network unicast or multicast address
<i>pingNb</i>	Number of ping slots per beacon period. This must be a power of 2 integer: $pingNb = 2^k$ where $1 \leq k \leq 7$
<i>pingPeriod</i>	Period of the device receiver wake-up expressed in number of slots: $pingPeriod = 2^{12} / pingNb$
<i>pingOffset</i>	Randomized offset computed at each beacon period start. Values can range from 0 to $(pingPeriod-1)$
<i>beaconTime</i>	The time carried in the field BCNPayload . Time of the immediately preceding beacon frame
<i>slotLen</i>	Length of a unit ping slot = 30 ms

6

7 At each beacon period the end-device and the server compute a new pseudo-random offset
8 to align the reception slots. An AES encryption with a fixed key of all zeros is used to
9 randomize:

10 $Key = 16 \times 0x00$

11 $Rand = aes128_encrypt(Key, beaconTime \parallel DevAddr \parallel pad16)$

12 $pingOffset = (Rand[0] + Rand[1] \times 256) \text{ modulo } pingPeriod$

13 The slots used for this beacon period will be:

14 $pingOffset + N \times pingPeriod$ with $N=[0:pingNb-1]$

15 The node therefore opens receive slots starting at :

First slot	Beacon_reserved + pingOffset x slotLen
Slot 2	Beacon_reserved + (pingOffset + pingPeriod) x slotLen
Slot 3	Beacon_reserved + (pingOffset + 2 x pingPeriod) x slotLen
...	...

16 If the end-device serves simultaneously a unicast and one or more multicast slots this
17 computation is performed multiple times at the beginning of a new beacon period. Once for
18 the unicast address (the node network address) and once for each multicast group address.

19 In the case where a multicast ping slot and a unicast ping slot collide and cannot be served
20 by the end-device receiver then the end-device should preferentially listen to the multicast
21 slot. If there is a collision between multicast reception slots the FPending bit of the previous
22 multicast frame can be used to set a preference.

23 The randomization scheme prevents a systematic collision between unicast and multicast
24 slots. If collisions happen during a beacon period then it is unlikely to occur again during the
25 next beacon period.

1 14 Class B MAC commands

2 All commands described in the Class A specification shall be implemented in Class B
3 devices. The Class B specification adds the following MAC commands.

4

CID	Command	Transmitted by		Short Description
		End-device	Gateway	
0x10	<i>PingSlotInfoReq</i>	x		Used by the end-device to communicate the ping unicast slot data rate and periodicity to the network server
0x10	<i>PingSlotInfoAns</i>		x	Used by the network to acknowledge a <i>PingInfoSlotReq</i> command
0x11	<i>PingSlotChannelReq</i>		x	Used by the network server to set the unicast ping channel of an end-device
0x11	<i>PingSlotFreqAns</i>	x		Used by the end-device to acknowledge a <i>PingSlotChannelReq</i> command
0x12	<i>BeaconTimingReq</i>	x		Used by end-device to request next beacon timing & channel to network
0x12	<i>BeaconTimingAns</i>		x	Used by network to answer a <i>BeaconTimingReq</i> uplink
0x13	<i>BeaconFreqReq</i>		x	Command used by the network server to modify the frequency at which the end-device expects to receive beacon broadcast
0x13	<i>BeaconFreqAns</i>	x		Used by the end-device to acknowledge a <i>BeaconFreqReq</i> command

5 14.1 PingSlotInfoReq

6 With the ***PingSlotInfoReq*** command an end-device informs the server of its unicast ping
7 slot periodicity and expected data rate. This command must only be used to inform the
8 server of the parameters of a UNICAST ping slot. A multicast slot is entirely defined by the
9 application and should not use this command.

10

Size (bytes)	1
PingSlotInfoReq Payload	Periodicity & data rate

11

Bit#	7	[6:4]	[3:0]
Periodicity & data rate	RFU	Periodicity	Data rate

12 The **Periodicity** subfield is an unsigned 3 bits integer encoding the ping slot period currently
13 used by the end-device using the following equation.

$$14 \quad pingSlotPeriod = 2^{Periodicity} \text{ in seconds}$$

- 15 • **Periodicity** = 0 means that the end-device opens a ping slot every second
- 16 • **Periodicity** = 7, every 128 seconds which is the maximum ping period supported by
17 the LoRaWAN Class B specification.

18 The **Data rate** subfield encodes the data rate at which the end-device expects any ping. This
19 uses the same encoding scheme that the ***LinkAdrReq*** command described in the Class A
20 specification.

- 1 The server needs to be aware of the end-device ping slot periodicity or expected data rate
 2 else Class B downlinks will not happen successfully. For that purpose the ***PingSlotInfoReq***
 3 MAC command **must be acknowledged** with a ***PingSlotInfoAns*** before the device can
 4 switch from class A to Class B. To change its ping slot scheduling or data rate a device
 5 should first revert to Class A , send the new parameters through a ***PingSlotInfoReq***
 6 command and get an acknowledge from the server through a ***PingSlotInfoAns*** . It can then
 7 switch back to Class B with the new parameters.
- 8 This command can be concatenated with any other MAC command in the **FHDRFOpt** field
 9 as described in the Class A specification frame format.

10 14.2 BeaconFreqReq

- 11 This command is sent by the server to the end-device to modify the frequency on which this
 12 end-device expects the beacon.

13

Octets	3
PingSlotChannelReqPayload	Frequency

14

- 15 The Frequency coding is identical to the ***NewChannelReq*** MAC command defined in the
 16 Class A.

- 17 **Frequency** is a 24bits unsigned integer. The actual beacon channel frequency in Hz is 100
 18 x frequ. This allows defining the beacon channel anywhere between 100 MHz to 1.67 GHz
 19 by 100 Hz step. The end-device has to check that the frequency is actually allowed by its
 20 radio hardware and return an error otherwise.

- 21 A valid non-zero Frequency will force the device to listen to the beacon on a fixed frequency
 22 channel even if the default behavior specifies a frequency hopping beacon (i.e US ISM
 23 band).

- 24 A value of 0 instructs the end-device to use the default beacon frequency plan as defined in
 25 the “Beacon physical layer” section. Where applicable the device resumes frequency
 26 hopping beacon search.

27 14.3 PingSlotChannelReq

- 28 This command is sent by the server to the end-device to modify the frequency on which this
 29 end-device expects the downlink pings.

30

Octets	3	1
PingSlotChannelReqPayload	Frequency	DrRange

31

- 32 The Frequency coding is identical to the ***NewChannelReq*** MAC command defined in the
 33 Class A.

- 34 **Frequency** is a 24bits unsigned integer. The actual ping channel frequency in Hz is 100 x
 35 frequ. This allows defining the ping channel anywhere between 100MHz to 1.67GHz by
 36 100Hz step. The end-device has to check that the frequency is actually allowed by its radio
 37 hardware and return an error otherwise.

- 38 A value of 0 instructs the end-device to use the default frequency plan.

1 **DrRange** is the data rate range allowed on this channel. This byte is split in two 4-bit
2 indexes.

Bits	7:4	3:0
DrRange	Max data rate	Min data rate

4
5 Following the convention defined in the LoRaWAN Regional Parameters document
6 [PARAMS], the “Min data rate” subfield designates the lowest data rate allowed on this
7 channel. For example 0 designates DR0 / 125 kHz in the EU physical layer. Similarly “Max
8 data rate” designates the highest data rate. For example in the EU spec, DrRange = 0x77
9 means that only 50 kbps GFSK is allowed on a channel and DrRange = 0x50 means that
10 DR0 / 125 kHz to DR5 / 125 kHz are supported.

11 Upon reception of this command the end-device answers with a **PingSlotFreqAns**
12 message. The MAC payload of this message contains the following information:

Size (bytes)	1
pingSlotFreqAns Payload	Status

14 The **Status** bits have the following meaning:

Bits	7:2	1	0
Status	RFU	Data rate range ok	Channel frequency ok

	Bit = 0	Bit = 1
Data rate range ok	The designated data rate range exceeds the ones currently defined for this end device, the previous range is kept	The data rate range is compatible with the possibilities of the end device
Channel frequency ok	The device cannot use this frequency, the previous ping frequency is kept	The device is able to use this frequency.

16 14.4 BeaconTimingReq

17 This command is sent by the end-device to request the next beacon timing and channel.
18 This MAC command has no payload. The **BeaconTimingReq** & **BeaconTimingAns**
19 mechanism is only meant to accelerate the initial beacon search to lower the end-device
20 energy requirement.

21 The network may answer only a limited number of requests per a given time period. An end-
22 device must not expect that **BeaconTimingReq** is answered immediately with a
23 **BeaconTimingAns**. Class A end-devices wanting to switch to Class B should not transmit
24 more than one **BeaconTimingReq** per hour.

25 End-devices requiring a fast beacon lock must implement an autonomous beacon finding
26 algorithm.

1 14.5 BeaconTimingAns

2 This command is sent by the network to answer a **BeaconInfoReq** request.

Size (bytes)	2	1
BeaconInfoReqPayload	Delay	Channel

3 The “**Delay**” field is a 16bits unsigned integer. If the remaining time between the end of the
 4 current downlink frame and the start of the next beacon frame is noted *RTime* then:

5 $30\text{ ms} \times (\mathbf{Delay}+1) > RTime \geq 30\text{ ms} \times \mathbf{Delay}$


6 In networks where the beacon uses alternatively several channels, the “**Channel**” field is the
 7 index of the beaoning channel on which the next beacon will be broadcasted. For networks
 8 where the beacon broadcast frequency is fixed then this field content is 0.
 9

1 15 Beaconing (Class B option)

2 15.1 Beacon physical layer

3 Besides relaying messages between end-devices and network servers, all gateways
 4 participate in providing a time-synchronization mechanisms by sending beacons at regular
 5 fixed intervals configurable per network (BEACON_INTERVAL). All beacons are transmitted
 6 in radio packet implicit mode, that is, without a LoRa physical header and with no CRC being
 7 appended by the radio.

8



9 The beacon Preamble begins with (a longer than default) 10 unmodulated symbols. This
 10 allows end-devices to implement a low power duty-cycled beacon search.

11 The beacon frame length is tightly coupled to the operation of the radio Physical layer.
 12 Therefore the actual frame length might change from one region implementation to another.
 13 The changing fields are highlighted in **Bold** in the following sections.

14 15.1.1 EU 863-870MHz ISM Band

15 The beacons are transmitted using the following settings

DR	3	Corresponds to SF9 spreading factor with 125 kHz BW
CR	1	Coding rate = 4/5
frequency	869.525 MHz	This is the recommended frequency allowing +27 dBm EIRP. Network operators may use a different frequency as long as ETSI compliance is achieved

16 The beacon frame content is:

Size (bytes)	3	4	1	7	2
BCNPayload	NetID	Time	CRC	GwSpecific	CRC

17 15.1.2 US 902-928MHz ISM Band

18 The beacons are transmitted using the following settings:

DR	10	Corresponds to SF10 spreading factor with 500kHz bw
CR	1	Coding rate = 4/5
frequencies	923.3 to 927.5MHz with 600kHz steps	Beaconing is performed on the same channel that normal downstream traffic as defined in the Class A specification

19 The downstream channel used for a given beacon is:

$$20 \text{ Channel} = \left[\text{floor} \left(\frac{\text{beacon_time}}{\text{beacon_period}} \right) \right] \text{ modulo } 8$$

21 • whereby beacon_time is the integer value of the 4 bytes “Time” field of the beacon
 22 frame

23 • whereby beacon_period is the periodicity of beacons , 128 seconds

24 • whereby $\text{floor}(x)$ designates rounding to the integer immediately inferior to x

25

Example: the first beacon will be transmitted on 923.3MHz , the second on 923.9MHz, the 9th beacon will be on 923.3MHz again.

Beacon channel nb	Frequency [MHz]
0	923.3
1	923.9
2	924.5
3	925.1
4	925.7
5	926.3
6	926.9
7	927.5

The beacon frame content is:

Size (bytes)	3	4	2	7	1	2
BCNPayload	NetID	Time	CRC	GwSpecific	RFU	CRC

15.2 Beacon frame content

The beacon payload **BCNPayload** consists of a network common part and a gateway-specific part.

Size (bytes)	3	4	1/2	7	0/1	2
BCNPayload	NetID	Time	CRC	GwSpecific	RFU	CRC

The network common part contains a network identifier **NetID** to uniquely identify the network for which the beacon is sent, and a timestamp **Time** in seconds since 00:00:00 [Coordinated Universal Time](#) (UTC), 1 January 1970. The integrity of the beacon's network common part is protected by an 8 or 16 bits CRC depending on PHY layer parameters. The CRC-16 is computed on the NetID+Time fields as defined in the IEEE 802.15.4-2003 section 7.2.1.8. When an 8 bits CRC is required then the 8 LSBs of the computed CRC-16 are used.

For example: This is a valid EU868 beacon frame:

AA BB CC | 00 00 02 CC | 7E | 00 | 01 20 00 | 00 81 03 | DE 55

Bytes are transmitted left to right. The corresponding field values are:

Field	NetID	Time	CRC	InfoDesc	lat	long	CRC
Value Hex	CCBBAA	CC020000	7E	0	002001	038100	55DE

The CRC-16 of the NetID+Time fields is 0xC87E but only the 8LSBs are used in that case

The seven LSB of the **NetID** are called **NwkID** and match the seven MSB of the short address of an end-device. Neighboring or overlapping networks **must have** different **NwkIDs**.

1 The gateway specific part provides additional information regarding the gateway sending a
 2 beacon and therefore may differ for each gateway. The RFU field when applicable (region
 3 specific) should be equal to 0. The optional part is protected by a CRC-16 computed on the
 4 GwSpecific+RFU fields. The CRC-16 definition is the same as for the mandatory part.

5 For example: This is a valid US900 beacon:

Field	NetID	Time	CRC	InfoDesc	lat	long	RFU	CRC
Value Hex	CCBBAA	CC020000	C87E	0	002001	038100	00	D450

6 Over the air the bytes are sent in the following order:

7 AA BB CC | 00 00 02 CC | 7E C8 | 00 | 01 20 00 | 00 81 03 | 00 | 50 D4

8 Listening and synchronizing to the network common part is sufficient to operate a stationary
 9 end-device in Class B mode. A mobile end-device should also demodulate the gateway
 10 specific part of the beacon to be able to signal to the network server whenever he is moving
 11 from one cell to another.

12 **Note:** As mentioned before, all gateways send their beacon at exactly
 13 the same point in time (i.e., time-synchronously) so that for network
 14 common part there are no visible on-air collisions for a listening end-
 15 device even if the end-device simultaneously receives beacons from
 16 several gateways. With respect to the gateway specific part, collision
 17 occurs but an end-device within the proximity of more than one
 18 gateway will still be able to decode the strongest beacon with high
 19 probability.

20 15.3 Beacon GwSpecific field format

21 The content of the **GwSpecific** field is as follow:

Size (bytes)	1	6
GwSpecific	InfoDesc	Info

22 The information descriptor **InfoDesc** describes how the information field **Info** shall be
 23 interpreted.

24

InfoDesc	Meaning
0	GPS coordinate of the gateway first antenna
1	GPS coordinate of the gateway second antenna
2	GPS coordinate of the gateway third antenna
3:127	RFU
128:255	Reserved for custom network specific broadcasts

25 For a single omnidirectional antenna gateway the **InfoDesc** value is 0 when broadcasting
 26 GPS coordinates. For a site featuring 3 sectorized antennas for example, the first antenna
 27 broadcasts the beacon with **InfoDesc** equals 0, the second antenna with **InfoDesc** field
 28 equals 1, etc ...

1 15.3.1 Gateway GPS coordinate: InfoDesc = 0, 1 or 2

2 For **InfoDesc** = 0, 1 or 2, the content of the **Info** field encodes the GPS coordinates of the
3 antenna broadcasting the beacon

Size (bytes)	3	3
Info	Lat	Lng

4 The latitude and longitude fields (**Lat** and **Lng**, respectively) encode the geographical
5 location of the gateway as follows:

- 6 • The north-south latitude is encoded using a signed 24 bit word where -2^{23}
7 corresponds to 90° south (the South Pole) and 2^{23} corresponds to 90° north (the
8 North Pole). The equator corresponds to 0.
- 9 • The east-west longitude is encoded using a signed 24 bit word where -
10 2^{23} corresponds to 180° west and 2^{23} corresponds to 180° east. The Greenwich
11 meridian corresponds to 0.

12 15.4 Beacons precise timing

13 The beacon is sent every 128 seconds starting at 00:00:00 Coordinated Universal Time
14 (UTC), 1 January 1970 plus **NwkID** plus TBeaconDelay. Therefore the beacon is sent at

$$15 \quad B_T = k * 128 + \mathbf{NwkID} + T\text{BeaconDelay}$$

16 seconds after 00:00:00 Coordinated Universal Time (UTC), 1 January 1970

17 whereby k is the smallest integer for which

$$18 \quad k * 128 + \mathbf{NwkID} > T$$

19 whereby

20 T = seconds since 00:00:00 Coordinated Universal Time (UTC), 1 January 1970.

21 **Note:** T is not (!) Unix time. Similar to GPS time and unlike Unix time,
22 T is strictly monotonically increasing and is not influenced by leap
23 seconds.

24 Whereby TBeaconDelay is a network specific delay in the [0:50] ms range.

25 TBeaconDelay may vary from one network to another and is meant to allow a slight
26 transmission delay of the gateways. TBeaconDelay must be the same for all gateways of a
27 given network. TBeaconDelay must be smaller than 50 ms. All end-devices ping slots use
28 the beacon transmission time as a timing reference, therefore the network server as to take
29 TBeaconDelay into account when scheduling the class B downlinks.
30

31 15.5 Network downlink route update requirements

32 When the network attempts to communicate with an end-device using a Class B downlink
33 slot, it transmits the downlink from the gateway which was closest to the end-device when
34 the last uplink was received. Therefore the network server needs to keep track of the rough
35 position of every Class B device.

36 Whenever a Class B device moves and changes cell, it needs to communicate with the
37 network server in order to update its downlink route. This update can be performed simply
38 by sending a “confirmed” or “unconfirmed” uplink, possibly without applicative payload.

39 The end-device has the choice between 2 basic strategies:

- 1 • Systematic periodic uplink: simplest method that doesn't require demodulation of the
2 "gateway specific" field of the beacon. Only applicable to slowly moving or stationery
3 end-devices. There are no requirements on those periodic uplinks.
- 4 • Uplink on cell change: The end-device demodulates the "gateway specific" field of
5 the beacon, detects that the ID of the gateway broadcasting the beacon it
6 demodulates has changed, and sends an uplink. In that case the device should
7 respect a pseudo random delay in the [0:120] seconds range between the beacon
8 demodulation and the uplink transmission. This is required to insure that the uplinks
9 of multiple Class B devices entering or leaving a cell during the same beacon period
10 will not systematically occur at the same time immediately after the beacon
11 broadcast.
- 12 Failure to report cell change will result in Class B downlink being temporary not operational.
13 The network server may have to wait for the next end-device uplink to transmit downlink
14 traffic.
15
16

1 16 Class B unicast & multicast downlink channel frequencies

2 16.1 EU 863-870MHz ISM Band

3 All unicast&multicastClass B downlinks use a single frequency channel defined by the
4 “**PingSlotChannelReq**” MAC command. The default frequency is 869.525MHz

5 16.2 US 902-928MHz ISM Band

6 By default Class B downlinks use a channel function of the Time field of the last beacon (see
7 Beacon Frame content) and the DevAddr.

$$\text{Class B downlink channel} = \left[\text{DevAddr} + \text{floor} \left(\frac{\text{Beacon_Time}}{\text{Beacon_period}} \right) \right] \text{ modulo } 8$$

- 8 • Whereby Beacon_Time is the 32 bit Time field of the current beacon period
- 9 • Beacon_period is the length of the beacon period (defined as 128sec in the
10 specification)
- 11 • Floor designates rounding to the immediately lower integer value
- 12 • DevAddr is the 32 bits network address of the device

13 Class B downlinks therefore hop across 8 channels in the ISM band and all Class B end-
14 devices are equally spread amongst the 8 downlink channels.

15 If the “**PingSlotChannelReq**” command with a valid non-zero argument is used to set the
16 Class B downlink frequency then all subsequent ping slots should be opened on this single
17 frequency independently of the last beacon frequency.

18 If the “**PingSlotChannelReq**” command with a zero argument is sent, the end-device
19 should resume the default frequency plan, id Class B ping slots hopping across 8 channels.

20 The underlying idea is to allow network operators to configure end-devices to use a single
21 proprietary dedicated frequency band for the Class B downlinks if available, and to keep as
22 much frequency diversity as possible when the ISM band is used.

23

1

CLASS C – CONTINUOUSLY LISTENING

1 **17 Class C: Continuously listening end-device**

2 The end-devices implanting the Class C option are used for applications that have sufficient
3 power available and thus do not need to minimize reception time.

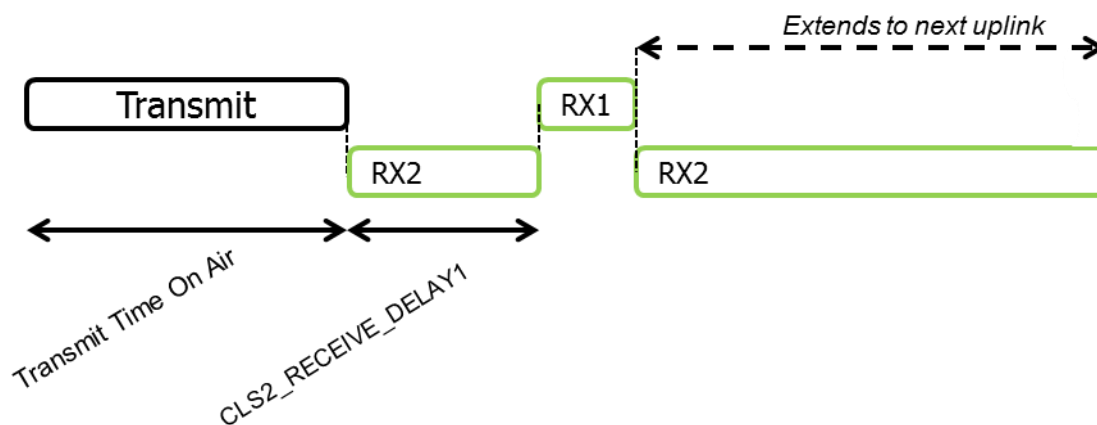
4 Class C end-devices cannot implement Class B option.

5 The Class C end-device will listen with RX2 windows parameters as often as possible. The
6 end-device listens on RX2 when it is not either (a) sending or (b) receiving on RX1,
7 according to Class A definition. To do so, it will open a short window on RX2 parameters
8 between the end of the uplink transmission and the beginning of the RX1 reception window
9 and it will switch to RX2 reception parameters as soon as the RX1 reception window is
10 closed; the RX2 reception window will remain open until the end-device has to send another
11 message.

12 **Note:** There is not specific message for a node to tell the server that it
13 is a Class C node. It is up to the application on server side to know that
14 it manages Class C nodes based on the contract passed during the
15 join procedure.

16 **17.1 Second receive window duration for Class C**

17 Class C devices implement the same two receive windows as Class A devices, but they do
18 not close RX2 window until they need to send again. Therefore they may receive a downlink
19 in the RX2 window at nearly any time, including downlinks sent for the purpose of MAC
20 command or ACK transmission. A short listening window on RX2 frequency and data rate is
21 also opened between the end of the transmission and the beginning of the RX1 receive
22 window.



23
24
25 **Figure 13: Class C end-device reception slot timing.**

26 **17.2 Class C Multicast downlinks**

27 Similarly to Class B, Class C devices may receive multicast downlink frames. The multicast
28 address and associated network session key and application session key must come from
29 the application layer. The same limitations apply for Class C multicast downlink frames:

- 30 • They are not allowed to carry MAC commands, neither in the **FOpt** field, nor in the
31 payload on port 0 because a multicast downlink does not have the same
32 authentication robustness as a unicast frame.

- 1 • The **ACK** and **ADRACKReq** bits must be zero. The **MType** field must carry the value
- 2 for Unconfirmed Data Down.
- 3 • The **FPending** bit indicates there is more multicast data to be sent. Given that a
- 4 Class C device keeps its receiver active most of the time, the **FPending** bit does not
- 5 trigger any specific behavior of the end-device.

1 **SUPPORT INFORMATION**

2 This sub-section is only a recommendation.

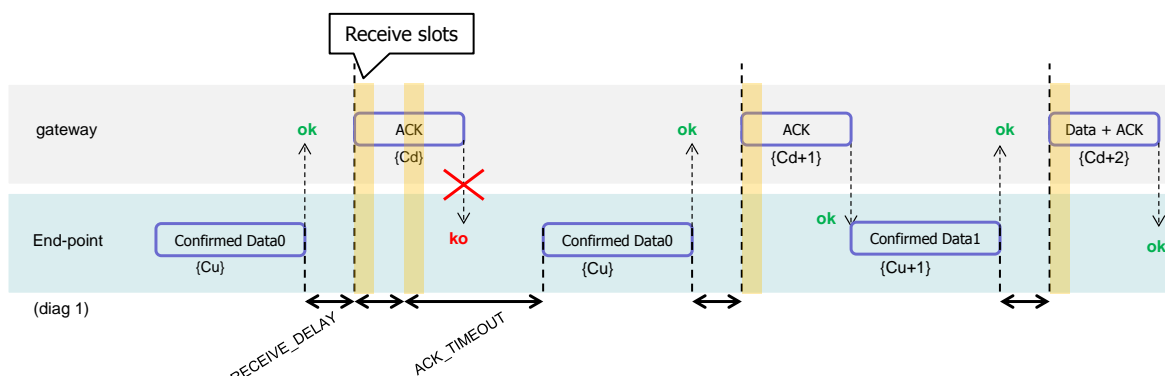
3

1 **18 Examples and Application Information**

2 Examples are illustrations of the LoRaWAN spec for information, but they are not part of the
3 formal specification.

4 **18.1 Uplink Timing Diagram for Confirmed Data Messages**

5 The following diagram illustrates the steps followed by an end-device trying to transmit two
6 confirmed data frames (Data0 and Data1):
7



8 **Figure 14: Uplink timing diagram for confirmed data messages**

9
10 The end-device first transmits a confirmed data frame containing the Data0 payload at an
11 arbitrary instant and on an arbitrary channel. The frame counter Cu is simply derived by
12 adding 1 to the previous uplink frame counter. The network receives the frame and
13 generates a downlink frame with the ACK bit set exactly RECEIVE_DELAY1 seconds later,
14 using the first receive window of the end-device. This downlink frame uses the same data
15 rate and the same channel as the Data0 uplink. The downlink frame counter Cd is also
16 derived by adding 1 to the last downlink towards that specific end-device. If there is no
17 downlink payload pending the network shall generate a frame without a payload. In this
18 example the frame carrying the ACK bit is not received.

19 If an end-device does not receive a frame with the ACK bit set in one of the two receive
20 windows immediately following the uplink transmission it may resend the same frame
21 with the same payload and frame counter again at least ACK_TIMEOUT seconds after the
22 second reception window. This resend must be done on another channel and must obey the
23 duty cycle limitation as any other normal transmission. If this time the end-device
24 receives the ACK downlink during its first receive window, as soon as the ACK frame is
25 demodulated, the end-device is free to transmit a new frame on a new channel.

26 The third ACK frame in this example also carries an application payload. A downlink frame
27 can carry any combination of ACK, MAC control commands and payload.

28 **18.2 Downlink Diagram for Confirmed Data Messages**

29 The following diagram illustrates the basic sequence of a “confirmed” downlink.
30
31

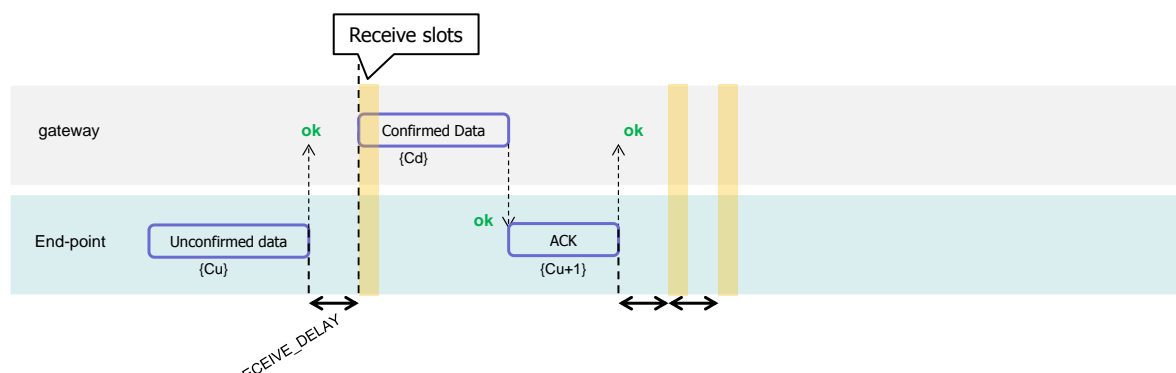


Figure 15: Downlink timing diagram for confirmed data messages

The frame exchange is initiated by the end-device transmitting an “unconfirmed” application payload or any other frame on channel A. The network uses the downlink receive window to transmit a “confirmed” data frame towards the end-device on the same channel A. Upon reception of this data frame requiring an acknowledgement, the end-device transmits a frame with the ACK bit set at its own discretion. This frame might also contain piggybacked data or MAC commands as its payload. This ACK uplink is treated like any standard uplink, and as such is transmitted on a random channel that might be different from channel A.

Note: To allow the end-devices to be as simple as possible and have keep as few states as possible it may transmit an explicit (possibly empty) acknowledgement data message immediately after the reception of a data message requiring an acknowledgment. Alternatively the end-device may defer the transmission of an acknowledgement to piggyback it with its next data message.

18.3 Downlink Timing for Frame-Pending Messages

The next diagram illustrates the use of the **frame pending** (FPending) bit on a downlink. The FPending bit can only be set on a downlink frame and informs the end-device that the network has several frames pending for him; the bit is ignored for all uplink frames.

If a frame with the FPending bit set requires an acknowledgement, the end-device shall do so as described before. If no acknowledgment is required, the end-device may send an empty data message to open additional receive windows at its own discretion, or wait until it has some data to transmit itself and open receive windows as usual.

Note: The FPending bit is independent to the acknowledgment scheme.

(*) F_P means 'frame pending' bit set

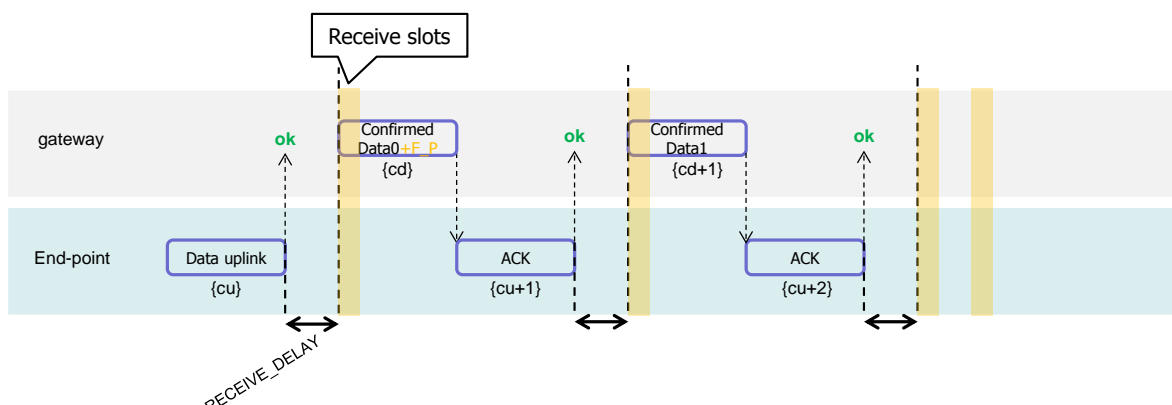
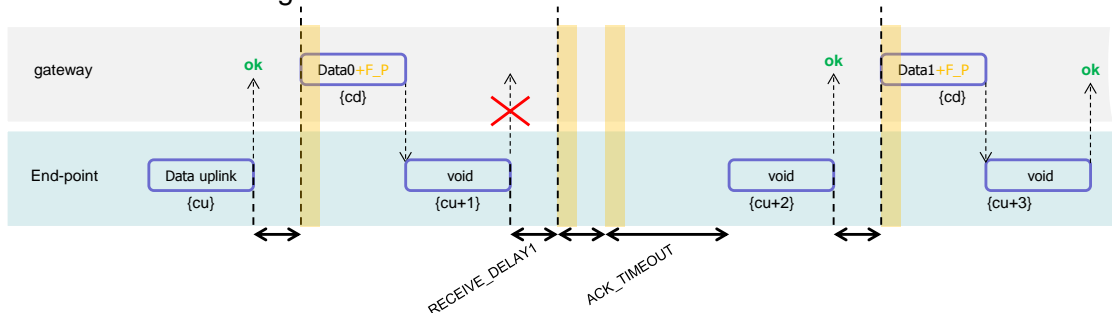


Figure 16: Downlink timing diagram for frame-pending messages, example 1

1 In this example the network has two confirmed data frames to transmit to the end-device.
 2 The frame exchange is initiated by the end-device via a normal “unconfirmed” uplink
 3 message on channel A. The network uses the first receive window to transmit the Data0 with
 4 the bit FPending set as a confirmed data message. The device acknowledges the reception
 5 of the frame by transmitting back an empty frame with the ACK bit set on a new channel B.
 6 RECEIVE_DELAY1 seconds later, the network transmits the second frame Data1 on
 7 channel B, again using a confirmed data message but with the FPending bit cleared. The
 8 end-device acknowledges on channel C.




9
10 **Figure 17: Downlink timing diagram for frame-pending messages, example 2**

11 In this example, the downlink frames are “unconfirmed” frames, the end-device does not
 12 need to send back and acknowledge. Receiving the Data0 unconfirmed frame with the
 13 FPending bit set the end-device sends an empty data frame. This first uplink is not received
 14 by the network. If no downlink is received during the two receive windows, the network has
 15 to wait for the next spontaneous uplink of the end-device to retry the transfer. The end-
 16 device can speed up the procedure by sending a new empty data frame.

17 **Note:** An acknowledgement is never sent twice.

18 The FPending bit, the ACK bit, and payload data can all be present in the same downlink.
 19 For example, the following frame exchange is perfectly valid.



21
22 **Figure 18: Downlink timing diagram for frame-pending messages, example 3**

23 The end-device sends a “confirmed data” uplink. The network can answer with a confirmed
 24 downlink containing Data + ACK + “Frame pending” then the exchange continues as
 25 previously described.

26 **18.4 Data-Rate Adaptation during Message Retransmissions**

27 When an end-device attempts the transmission of a “confirmed” frame toward the network it
 28 expects to receive an acknowledgement in one of the subsequent reception slot. In the
 29 absence of the acknowledgement it will try to re-transmit the same data again. This re-

1 transmission happens on a new frequency channel, but can also happen at a different data
 2 rate (preferable lower) than the previous one. It is strongly recommended to adopt the
 3 following re-transmission strategy.

4 The first transmission of the “confirmed” frame happens with a data rate DR.
 5

Transmission nb	Data Rate
1 (first)	DR
2	DR
3	$\max(\text{DR}-1,0)$
4	$\max(\text{DR}-1,0)$
5	$\max(\text{DR}-2,0)$
6	$\max(\text{DR}-2,0)$
7	$\max(\text{DR}-3,0)$
8	$\max(\text{DR}-3,0)$

6 The Data Rate $\max(a,b)$ stands for maximum of a and b values.

7 If after a recommended 8 transmissions, the frame has not been acknowledged the MAC
 8 layer should return an error code to the application layer.

9 **Note:** For each re-transmission, the frequency channel is selected
 10 randomly as for normal transmissions.

11 Any further transmission uses the last data rate used.

12 For example if an end-device sends a “confirmed” frame first using DR5 and has to
 13 retransmit 3 times (twice at DR5 then twice at DR4), the next frame transmitted will use DR4

14 Other example, if an end-device sends a “confirmed” frame first using DR5 and does not
 15 receive an acknowledge after 8 transmissions (2 at DR5, 2 at DR4, ... , 2 at DR2), and the
 16 application of this end-device re-initiates a “confirmed” transmission a little later, the first two
 17 transmission will be tentatively at DR2, then switch to DR1, then to DR0.

1 **19 Recommendation on contract to be provided to the network**
2 **server by the end-device provider at the time of provisioning**

3 Configuration data related to the end-device and its characteristics must be known by the
4 network server at the time of provisioning. –This provisioned data is called the “contract”.
5 This contract cannot be provided by the end-device and must be supplied by the end-device
6 provider using another channel (out-of-band communication).

7 This end-device contract is stored in the network server. It can be used by the application
8 server and the network controller to adapt the algorithms.

9 This data will include:

- 10 • End-device specific radio parameters (device frequency range, device maximal
11 output power, device communication settings - RECEIVE_DELAY1,
12 RECEIVE_DELAY2)
- 13 • Application type (Alarm, Metering, Asset Tracking, Supervision, Network Control)

1 20 Recommendation on finding the locally used channels

2 End-devices that can be activated in territories that are using different frequencies for
3 LoRaWAN will have to identify what frequencies are supported for join message at their
4 current location before they send any message. The following methods are proposed:

- 5 • A GPS enabled end-device can use its GPS location to identify which frequency
6 band to use.
- 7 • End-device can search for a beacon and use its frequency to identify its region
- 8 • End-device can search for a beacon and if this one is sending the antenna GPS
9 coordinate, it can use this to identify its region
- 10 • End-device can search for a beacon and if this one is sending a list of join
11 frequencies, it can use this to send its join message

1 21 Revisions

2 21.1 Revision 1.0

- 3 • Approved version of LoRaWAN1.0

4 21.2 Revision 1.0.1

- 5 • Clarified the RX window start time definition
- 6 • Corrected the maximum payload size for DR2 in the NA section
- 7 • Corrected the typo on the downlink data rate range in 7.2.2
- 8 • Introduced a requirement for using coding rate 4/5 in 7.2.2 to guarantee a maximum
- 9 time on air < 400mSec
- 10 • Corrected the JoinAccept MIC calculation in 6.2.5
- 11 • Clarified the NbRep field and renamed it to NbTrans in 5.2
- 12 • Removed the possibility to not encrypt the Applicative payload in the MAC layer ,
- 13 removed the paragraph 4.3.3.2. If further security is required by the application , the
- 14 payload will be encrypted, using any method, at the application layer then re-
- 15 encrypted at the MAC layer using the specified default LoRaWAN encryption
- 16 • Corrected FHDR field size typo
- 17 • Corrected the channels impacted by ChMask when chMaskCntl equals 6 or 7 in
- 18 7.2.5
- 19 • Clarified 6.2.5 sentence describing the RX1 slot data rate offset in the JoinResp
- 20 message
- 21 • Removed the second half of the DRoffset table in 7.2.7 , as DR>4 will never be used
- 22 for uplinks by definition
- 23 • Removed explicit duty cycle limitation implementation in the EU868Mhz ISM band
- 24 (chapter7.1)
- 25 • Made the RXtimingSetupAns and RXParamSetupAns sticky MAC commands to
- 26 avoid end-device's hidden state problem. (in 5.4 and 5.7)
- 27 • Added a frequency plan for the Chinese 470-510MHz metering band
- 28 • Added a frequency plan for the Australian 915-928MHz ISM band

29 21.3 Revision 1.0.2

- 30 • Extracted section 7 “Physical layer” that will now be a separated document
- 31 “LoRaWAN Regional Parameters”
- 32 • corrected the ADR_backoff sequence description (ADR_ACK_LIMT was written
- 33 instead of ADR_ACK_DELAY) paragraph 4.3.1.1
- 34 • Corrected a formatting issue in the title of section 18.2 (previously section 19.2 in the
- 35 1.0.1 version)
- 36 • Added the DIChannelRec MAC command, this command is used to modify the
- 37 frequency at which an end-device expects a downlink.
- 38 • Added the Tx ParamSetupRec MAC command. This command enables to remotely
- 39 modify the maximum TX dwell time and the maximum radio transmit power of a
- 40 device in certain regions
- 41 • Added the ability for the end-device to process several ADRreq commands in a
- 42 single block in 5.2
- 43 • Clarified AppKey definition
- 44

1 22 Glossary

2		
3	ADR	Adaptive Data Rate
4	AES	Advanced Encryption Standard
5	AFA	Adaptive Frequency Agility
6	AR	Acknowledgement Request
7	CBC	Cipher Block Chaining
8	CMAC	Cipher-based Message Authentication Code
9	CR	Coding Rate
10	CRC	Cyclic Redundancy Check
11	DR	Data Rate
12	ECB	Electronic Code Book
13	ETSI	European Telecommunications Standards Institute
14	EIRP	Equivalent Isotropically Radiated Power
15	FSK	Frequency Shift Keying modulation technique
16	GPRS	General Packet Radio Service
17	HAL	Hardware Abstraction Layer
18	IP	Internet Protocol
19	LBT	Listen Before Talk
20	LoRa™	Long Range modulation technique
21	LoRaWAN™	Long Range Network protocol
22	MAC	Medium Access Control
23	MIC	Message Integrity Code
24	RF	Radio Frequency
25	RFU	Reserved for Future Usage
26	Rx	Receiver
27	RSSI	Received Signal Strength Indicator
28	SF	Spreading Factor
29	SNR	Signal Noise Ratio
30	SPI	Serial Peripheral Interface
31	SSL	Secure Socket Layer
32	Tx	Transmitter
33	USB	Universal Serial Bus
34		

1 **23 Bibliography**

2 **23.1 References**

- 3 [IEEE802154]: IEEE Standard for Local and Metropolitan Area Networks—Part 15.4: Low-
4 Rate Wireless Personal Area Networks (LR-WPANs), IEEE Std 802.15.4TM-2011 (Revision
5 of IEEE Std 802.15.4-2006), September 2011.
- 6 [RFC4493]: The AES-CMAC Algorithm, June 2006.
- 7 [PARAMS]: LoRaWAN Regional Parameters, the LoRa Alliance, October 2016 (and future
8 versions).

1 24 NOTICE OF USE AND DISCLOSURE

2 Copyright © LoRa Alliance, Inc. (2015). All Rights Reserved.

3 The information within this document is the property of the LoRa Alliance (“The Alliance”) and its use and
4 disclosure are subject to LoRa Alliance Corporate Bylaws, Intellectual Property Rights (IPR) Policy and
5 Membership Agreements.

6 Elements of LoRa Alliance specifications may be subject to third party intellectual property rights, including
7 without limitation, patent, copyright or trademark rights (such a third party may or may not be a member of LoRa
8 Alliance). The Alliance is not responsible and shall not be held responsible in any manner for identifying or failing
9 to identify any or all such third party intellectual property rights.

10 This document and the information contained herein are provided on an “AS IS” basis and THE ALLIANCE
11 DISCLAIMS ALL WARRANTIES EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO (A) ANY
12 WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OF THIRD
13 PARTIES (INCLUDING WITHOUT LIMITATION ANY INTELLECTUAL PROPERTY RIGHTS INCLUDING
14 PATENT, COPYRIGHT OR TRADEMARK RIGHTS) OR (B) ANY IMPLIED WARRANTIES OF
15 MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE OR NONINFRINGEMENT.

16 IN NO EVENT WILL THE ALLIANCE BE LIABLE FOR ANY LOSS OF PROFITS, LOSS OF BUSINESS, LOSS
17 OF USE OF DATA, INTERRUPTION OF BUSINESS, OR FOR ANY OTHER DIRECT, INDIRECT, SPECIAL OR
18 EXEMPLARY, INCIDENTAL, PUNITIVE OR CONSEQUENTIAL DAMAGES OF ANY KIND, IN CONTRACT OR
19 IN TORT, IN CONNECTION WITH THIS DOCUMENT OR THE INFORMATION CONTAINED HEREIN, EVEN IF
20 ADVISED OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

21 The above notice and this paragraph must be included on all copies of this document that are made.

22 LoRa Alliance, Inc.

23 2400 Camino Ramon, Suite 375

24 San Ramon, CA 94583

25 *Note: All Company, brand and product names may be trademarks that are the sole property of their respective*
26 *owners.*