TS001-1.0.4 LoRaWAN® L2 1.0.4 Specification

NOTICE OF USE AND DISCLOSURE

LoRa Alliance®
5177 Brandin Court
Fremont, CA 94538
United States

LoRaWAN® L2 1.0.4 Specification (TS001-1.0.4)

Authored by the LoRa Alliance Technical Committee

Technical Committee Chair and Vice-Chair:
A.YEGIN (Actility), O.SELLER (Semtech)

Editors:
T.KRAMP (Semtech), O.SELLER (Semtech)

Contributors (in alphabetical order):
A.BERTOLAUD (Gemalto), I.CALABRESE (A2A Smart City), J.CATALANO (Kerlink), J.DELCLEF (ST Microelectronics), V.DELPORT (Microchip Technology), P.DUFFY (Cisco), F.DYDUCH (Bouygues Telecom), T.EIRICH (Semtech), L.FERREIRA (Orange), Y.GAUDIN (Kerlink), S.GHAROUT (Orange), O.HERSENT (Actility), A.KASTTET (Birdz), D.KJENDAL (Senet), V.KLEBAN (Everynet), J.KNAPP (Semtech), T.KRAMP (Semtech), M.KUYPER (Semtech), P.KWOK (Objenious), M.LEGOURIEREC (Sagemcom), C.LEVASSEUR (Bouygues Telecom), M.LUIS (Semtech), M.PAULIAC (Gemalto), P.PIETRI (Orbiwise), O.SELLER (Semtech), D.SMITH (MultiTech), N.SORNIN (Semtech), R.SOSS (Actility), J.STOKKING (The Things Network), T.TASHIRO (M2B Communications), D.THOLL (Tektelic), P.THOMSEN (Orbiwise), A.YEGIN (Actility)

Version: 1.0.4
Date: October 2020
Status: Released

Contents

Tables

Figures

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 are typically laid out in a star-of-stars topology in which gateways[1] relay transmissions between end-devices and a central Network Server at the backend. Gateways are connected to a Network Server via standard IP connections, whereas end-devices use single-hop radio-frequency (RF) communication to one or many gateways. All communication is generally bi-directional, although uplink communication from an end-device to a Network Server is expected to be the predominant traffic.

Communication between end-devices and gateways is distributed over different frequency channels and data rates. Selecting the data rate is a tradeoff between communication range and transmission duration; communications with different LoRa data rates do not interfere with each other. To maximize both the battery life of end-devices and the overall network capacity, the LoRaWAN network infrastructure MAY manage the data rate and RF transmit power for each end-device individually by means of an adaptive data rate (ADR) scheme.

An end-device may transmit on any channel available at any time using any available data rate, as long as the following rules are observed:

- The end-device changes channels in a pseudo-random fashion for every transmission. The resulting frequency diversity makes the system more robust to interference.
- The end-device pseudo-randomly changes its transmit periodicity to prevent systematic synchronization of populations of end-device transmissions.
- The end-device complies with all local regulations governing its behavior in the band and sub-bands in which it is currently operating including, but not limited to, duty-cycle and dwell-time (transmit-duration) limitations.

All LoRaWAN end-devices SHALL implement at least Class A functionality as described in this document. In addition, they MAY implement Class B and/or Class C as also described in this document.

End-devices implementing Class B are referred to as Class B-capable. When operating in Class B, end-devices are referred to as Class B-enabled. Transition from Class B-disabled to Class B-enabled is called switching to Class B**.**

End-devices implementing Class C are referred to as Class C-capable. When operating in Class C, end-devices are referred to as Class C-enabled. Transition from Class C-disabled to Class C-enabled is called switching to Class C**.**

In all cases, end-devices remain compatible with Class A.

[1] Gateways are also known as concentrators, routers, access points, or base stations.

1.1 Conventions

The keywords "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

The tables in this document are normative. The figures in this document are informative. The notes in this document are informative.

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

In this document,

- The octet order for all multi-octet fields SHALL be little endian.
- EUI are 8-octet fields and SHALL be transmitted as little endian.
- By default, `RFU` bits are Reserved for Future Use and SHALL be set to 0 by the transmitter of the frame and SHALL be silently ignored by the receiver.

2 Introduction to LoRaWAN Options

LoRa® is a wireless modulation for long-range, low-power, low-data-rate applications developed by Semtech. End-devices implementing more than Class A are generally called "higher Class end-devices" in this document.

2.1 LoRaWAN Classes

A LoRaWAN network distinguishes between a basic LoRaWAN (called Class A) and optional features (Class B, Class C …) as shown in Figure 1.

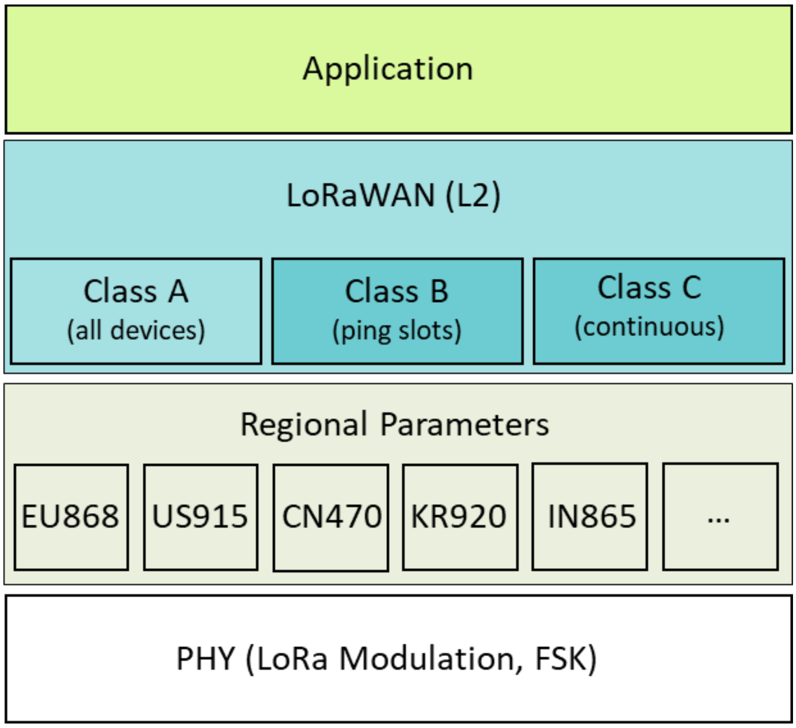

Figure 1: LoRaWAN Classes

- **Bi-directional end-devices (Class A):** Class A end-devices allow 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 protocol). Class A operation is the lowest-power end-device system for applications that require only 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 uplink initiated by the end-device.
- **Bi-directional end-devices with scheduled receive slots (Class B):** Class B-capable end-devices allow more receive slots. In addition to Class A receive windows, Class B-enabled end-devices open extra receive windows at scheduled times. In order for the end-device to open its receive windows at a scheduled time, it receives a time-synchronized beacon from the gateway. This allows the Network Server to know when the end-device is listening.

- **Bi-directional end-devices with maximal receive slots (Class C):** Class C-capable end-devices allow nearly continuously open receive windows, which are closed only when transmitting. Class C-enabled end-devices use more power to operate than Class A or Class B-enabled, but they feature the lowest latency for communication between servers and end-devices.

CLASS A – ALL END-DEVICES

All LoRaWAN end-devices SHALL implement all Class A features not explicitly marked optional.

> **Note:** Physical packet format, MAC frame format and other parts of this specification that are common to both end-devices of Class A and higher Classes are described only in the LoRaWAN Class A specification in order to avoid redundancy.

3 Physical Packet Formats

The physical layers used by LoRaWAN are defined in [RP002].

The LoRaWAN terminology distinguishes between uplink and downlink frames.

3.1 Uplink Packets

Uplink packets are sent by end-devices to a Network Server relayed by one or many gateways.

3.2 Downlink Packets

Downlink packets are sent by a Network Server to only one end-device and are transmitted by a Network Server through one or more gateways.[2]

3.3 Receive Windows

Following each uplink transmission, the end-device SHALL open one or two receive windows (RX1 and RX2); if no packet destined for the end-device is received in RX1, it SHALL open RX2. The receive windows start times are defined using the end of the transmission as a reference, see Figure 2.

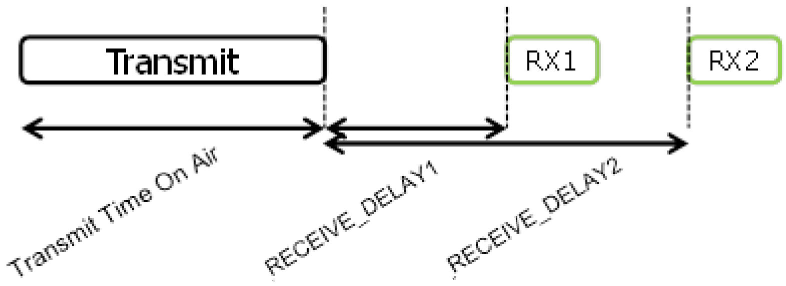

Figure 2: End-device receive-slot timing

3.3.1 Receiver activity during receive windows

If a preamble is detected during one of the receive windows, the radio receiver SHOULD stay active until the downlink frame is demodulated. If a frame was detected and subsequently demodulated during the first receive window, and the frame was intended for this end-device after address and MIC (message integrity code) checks, the end-device SHALL NOT open the second receive window.

[2] This specification does not describe the transmission of multicast frames from a Network Server to many end-devices.

3.3.2 First receive-window channel, data rate, and start

The first receive window RX1 uses a frequency that is a function of the uplink frequency and a data rate that is a function of the uplink data rate. RX1 SHALL be open no longer than RECEIVE_DELAY1 seconds after the end of the uplink modulation.[3] The relationship between uplink and RX1 data rates is region-specific and detailed in the "LoRaWAN Regional Parameters" [RP002] document.

3.3.3 Second receive window channel, data rate, and start

The second receive window RX2, if opened (see Section 3.3.1), uses a fixed configurable frequency and data rate, and SHALL be open no longer than RECEIVE_DELAY2 seconds after the end of the uplink modulation.[3] The frequency and data rate can be modified by MAC commands (see Section 5). The default frequency and data rate are region-specific and detailed in "LoRaWAN Regional Parameters" [RP002].

3.3.4 Receive window duration

The duration of a receive window SHALL be at least the time required by the end-device's radio transceiver to detect a downlink preamble starting at RECEIVE_DELAY1 or RECEIVE_DELAY2 after the end of the uplink modulation.

> **Note:** The end-device receive window duration has to accommodate the maximum potential imprecision of the end-device's clock. The delay between the end of the uplink and the start of the receive windows can be changed from 1s to 15s for RX1 (and thereby from 2s to 16s for RX2) using the OTAA (over-the-air activation) Join-Accept frame or the ***RXTimingSetupReq*** MAC command. Therefore, this delay must be accommodated when computing the maximum clock imprecision.
>
> **Example:** A 30 ppm XTAL frequency error translates to ±30 µs after 1 s and ±450 µs after 15 s.

3.3.5 Network transmitting to an end-device

If a Network Server intends to transmit a downlink to an end-device, it SHALL initiate the transmission of the downlink frame precisely at the beginning of one of those two receive windows. Such a downlink is referred to as a Class A downlink. The end-device SHALL open a Class A RX1 receive window. If no frame intended for this end-device was received during the RX1 receive window, the end-device SHALL open an RX2 receive window at the specified timing, even if this interrupts the reception of a transmission using Class B or Class C downlink timing and receive parameters.

3.3.6 Important notice regarding receive windows

An end-device SHALL NOT transmit another uplink packet before it has either received a downlink packet in the first or second receive window related to the previous transmission or if the second receive window related to the previous transmission has expired.

[3] RECEIVE_DELAY1 and RECEIVE_DELAY2 are described in Section 6.

3.3.7 Receiving or transmitting other protocols

The end-device MAY listen for or transmit other protocols or perform arbitrary processing between LoRaWAN transmission and reception windows, as long as the end-device remains compatible with local regulations and compliant with the LoRaWAN specification.

4 MAC Frame Formats

All LoRaWAN uplink and downlink packets carry a PHY payload (`PHYPayload`) starting with a single-octet MAC header (`MHDR`), followed by a MAC payload (`MACPayload`), and ending with a 4-octet message integrity code (`MIC`).

PHYPayload:

MHDR	MACPayload	MIC

or

MHDR	Join-Request[4]	MIC

or

MHDR	Join-Accept[4]	MIC

PHYPayload structure

MACPayload:

FHDR	FPort	FRMPayload

MACPayload structure

FHDR:

DevAddr	FCtrl	FCnt	FOpts

Frame header structure

Table 1: LoRaWAN frame format elements

4.1 PHY Payload (PHYPayload)

Size (octets)	1	7..*M*	4
PHYPayload	MHDR	MACPayload	MIC

Table 2: PHYPayload format

The maximum length *M* of the `MACPayload` field is region- and data-rate-specific and is specified in [RP002]. Neither the end-device nor the Network SHALL send a frame containing a `MACPayload` greater than the specified maximum length *M* over the data rate used to transmit the frame. Any frame received by an end-device or a Network Server containing a `MACPayload` greater than the specified maximum length *M* over the data rate used to receive the frame SHALL be silently discarded.

[4] Cf. Section 6.2.4

4.2 MAC Header (MHDR field)

Bits	[7:5]	[4:2]	[1:0]
MHDR	FType	RFU	Major

Table 3: MHDR format

The MAC header specifies the frame type (FType) and the Major version (Major) of the frame format of the LoRaWAN layer specification according to which the frame has been encoded.

4.2.1 Frame types (FType bit field)

The LoRaWAN distinguishes among six different MAC frame types: Join-Request, Join-Accept, unconfirmed data uplink / downlink, and confirmed data uplink / downlink.

FType	Description
000	Join-Request
001	Join-Accept
010	unconfirmed data uplink
011	unconfirmed data downlink
100	confirmed data uplink
101	confirmed data downlink
110	RFU
111	Proprietary

Table 4: MAC frame types

4.2.1.1 Join-Request and Join-Accept frames

Join-Request and Join-Accept frames are used by the over-the-air activation procedure described in Section 6.2.

4.2.1.2 Data frames

Data frames transfer MAC commands and application data, which can be combined into a single frame. A **confirmed data frame** SHALL be acknowledged by the receiver, whereas an **unconfirmed data frame** SHALL NOT be acknowledged.[5] **Proprietary frames** MAY be used to implement non-standard formats that are not interoperable with standard frames but SHALL be used only among end-devices that have a common understanding of the proprietary extensions.

Frame integrity is ensured in different ways for different frame types and is described per frame type below.

[5] A detailed timing diagram of the acknowledge mechanism is shown in Section 16.

4.2.2 Major data frame version (Major bit field)

Major	Description
00	LoRaWAN R1
01..11	RFU

Table 5: Major list

The Major version specifies the format of the frames exchanged in the Join procedure (see Section 6.2.4) and the first four octets of the MACPayload as described in Section 4. For each Major version, end-devices MAY implement different Minor versions of the frame format. The Minor version used by an end-device SHALL be made known to the Network Server beforehand using out-of-band communication (e.g., as part of the end-device personalization information).

4.3 MAC Payload of Data Frames (MACPayload)

The MAC payload of data frames contains a frame header (FHDR) followed by an OPTIONAL port field (FPort) and an OPTIONAL frame payload field (FRMPayload).

A frame with a valid FHDR, no FOpts (FOptsLen=0), no FPort and no FRMPayload is a valid frame.

4.3.1 Frame header (FHDR)

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

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

Table 6: FHDR format

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

Bits	7	6	5	4	[3..0]
FCtrl	ADR	RFU	ACK	FPending	FOptsLen

Table 7: FCtrl downlink frames format

For uplink frames, the `FCtrl` content of the frame header is

Bits	7	6	5	4	[3..0]
FCtrl	ADR	ADRACKReq	ACK	ClassB	FOptsLen

Table 8: `FCtrl` uplink frame format

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

LoRaWAN allows end-devices to use any of the possible data rates and transmit (TX) power individually. This feature is used by Network Servers to adapt and optimize the number of retransmissions, the data rate, and the TX power of end-devices. This is referred to as the adaptive data rate (ADR) and, when it is enabled, end-devices will be optimized to use the fastest data rate and minimum TX power possible.

ADR control may not be possible when radio channel attenuation changes rapidly and/or continuously. When the Network is unable to control the data rate of an end-device, the end-device's application layer SHOULD control it. It is RECOMMENDED that a variety of different data rates be used in this case. The application layer SHOULD always try to minimize the aggregated airtime, given the network conditions.

If the uplink `ADR` bit is set, the Network may control the number of retransmissions, the data rate, and the TX power of the end-device through the appropriate MAC commands. If the `ADR` bit is unset, the Network Server SHALL accept that the end-device MAY not comply with any attempt to control the number of retransmissions, the data rate, or the TX power of the end-device regardless of the received signal quality. The Network MAY still send commands to inform the end-device of the recommended configuration using the ***LinkADRReq*** command. An end-device SHOULD accept the channel mask controls present in ***LinkADRReq***, even when the `ADR` bit is not set. The end-device SHALL respond to all ***LinkADRReq*** commands with a ***LinkADRAns*** indicating which command elements were accepted and which were rejected. This behavior differs from when the uplink `ADR` bit is set, in which case the end-device accepts or rejects the entire command.

> **Note:** A Network Server may not infer any actual end-device state in the case where the uplink `ADR` bit is not set, regardless of the state of the individual `Status` bits of ***LinkADRAns***. These are provided for offline debugging.

When the downlink `ADR` bit is set, it informs the end-device that the Network Server is able to send ADR commands. The end-device MAY set/unset the uplink `ADR` bit independently.

When the downlink `ADR` bit is unset, it signals the end-device that, owing to rapid changes of the radio channel, the Network temporarily cannot estimate the best data rate. In that case, the end-device has the choice to

- Unset the `ADR` uplink bit and control its uplink data rate, TX power and channel plan following its own strategy. This SHOULD be the typical strategy for a mobile end-device, or
- Ignore it (keep the uplink `ADR` bit set) and apply the normal ADR backoff algorithm in the absence of downlinks. This SHOULD be the typical strategy for a stationary end-device.

The ADR bit MAY be set and unset on demand by the end-device or a Network Server. 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 immobile most of the time. Depending on its state of mobility, an end-device can request that the Network optimize its data rate using ADR.

Default TX power is the maximum transmit power allowed for an end-device, considering its capabilities and any applicable regional regulatory constraints. End-devices SHALL use this power level until the Network attempts to change it using the ***LinkADRReq*** MAC command, or if the end-device has unset the ADR bit.

The default data rate is the minimum data rate allowed for an end-device, considering its capabilities and any applicable regional regulatory constraints. An end-device using Activation By Personalization (ABP, see section 6), with the ADR bit set, SHALL use this data rate until the Network requests a higher data rate through the ***LinkADRReq*** MAC command. End-devices which use Over The Air Activation follow a Join Procedure (see section 6) which determines the initial uplink data rate.

If an end-device wishes to check for connectivity loss or if an end-device whose data rate is optimized by the Network uses a data rate higher than its default data rate or a TX power lower than its default, the end-device SHALL periodically validate whether the Network is still receiving the uplink frames. Each time the uplink frame counter is incremented (for each new uplink frame, because repeated transmissions do not increment the frame counter), the end-device SHALL increment an ADRACKCnt counter. After ADR_ACK_LIMIT uplinks (ADRACKCnt ≥ ADR_ACK_LIMIT) without receiving a Class A downlink response, the end-device SHALL set the ADR acknowledgment request bit (ADRACKReq) on uplink transmissions. The Network is REQUIRED to respond with a class A downlink frame within the next ADR_ACK_DELAY frames. A Class A downlink frame received following an uplink frame SHALL reset the ADRACKCnt counter. Upon receipt of any Class A downlink, the end-device SHALL clear the ADRACKReq bit. The downlink ACK bit does not need to be set because any Class A downlink frame received by the end-device indicates that the Network has received uplinks from this end-device. If no Class A downlink frame is received within the next ADR_ACK_DELAY uplinks (i.e., after a total of ADR_ACK_LIMIT + ADR_ACK_DELAY transmitted frames), the end-device SHALL try to regain connectivity by first setting the TX power to the default power, then switching to the next lower data rate that provides a longer radio range. The end-device SHALL further lower its data rate step by step every time ADR_ACK_DELAY uplink frames are transmitted. Once the end-device has reached the default data rate, and transmitted for ADR_ACK_DELAY uplinks with ADRACKReq=1 without receiving a downlink, it SHALL re-enable all default uplink frequency channels and reset NbTrans to its default value of 1. Furthermore, if at any point during the backoff the resulting configuration results in an invalid combination of TX power, data rate or channel mask, the end-device SHALL immediately re-enable all default channels and use the maximum TX power permissible for and available to this end-device.

> **Note:** Other configurations of the end-device by the Network are not modified during ADR backoff. Specifically, configurations that affect downlink connectivity (controlled by ***RXParamsSetupReq***, ***DlChannelReq***, ***RXTimingSetupReq***, and ***TXParamSetupReq***) are not modified during ADR backoff.

For fixed channel plan regions (US, AU, CN, etc.), end-devices SHALL enable all channels. For dynamic channel plan regions (EU, IN, AS, etc.), end-devices SHALL enable the region's default channels and make no change to the configuration of the dynamically configured channels.

> **Note:** Not requiring an immediate response to an ADR acknowledgment request provides flexibility to the Network to schedule its downlinks in an optimal manner.

Table 9 provides an example of a data rate backoff sequence, assuming ADR_ACK_LIMIT is set to 64 and ADR_ACK_DELAY is equal to 32.

ADRACKCnt	ADRACKReq	Data Rate	TX Power	NbTrans	Channel Mask
0 to 63	0	DR1	Max –9 dBm	3	Normal operations channel mask
64 to 95	1	No change	No change	No change	No change
96 to 127	1	No change	Default	No change	No change
128 to 159	1	DR0 (Default)	Default	No change	No change
≥ 160	1	DR0 (Default)	Default	1	For dynamic channel plans: re-enable default channels. For fixed channel plans: All channels enabled

Table 9: Example of a data rate backoff sequence

4.3.1.2 Frame acknowledge bit and acknowledgment procedure (ACK in FCtrl)

When receiving a confirmed data frame, the receiver responds with a data frame that has the acknowledgment bit (ACK) set. If the Network receives such a confirmed frame it SHOULD send an acknowledgment. If the Network sends an acknowledgement it SHALL send it using one of the Class A receive windows opened by the end-device after the send operation. If an end-device receives such a confirmed frame in one of its Class A receive windows, it SHALL transmit an acknowledgment with its next uplink. If an end-device receives such a confirmed frame outside of its Class A receive windows, i.e. in a Class B ping slot (see section 9) or in RXC (see section 15), it SHOULD send an acknowledgement.

Acknowledgments SHALL only be sent in response to the latest frame received and SHALL NOT be transmitted more than NbTrans times. The Network SHALL only send an acknowledgment in the Class A receive windows (RX1/RX2) of the confirmed uplink that requested it.

> **Note:** To allow an end-device to be as simple as possible and have as few states as possible, it may transmit an explicit (possibly empty) acknowledgment data frame immediately after receiving a data frame requiring a confirmation. Alternatively, an end-device may defer the transmission of an acknowledgment to piggyback it with its next data frame.

4.3.1.3 Retransmission procedure

Downlink frames

A downlink confirmed or unconfirmed frame SHALL NOT be retransmitted using the same frame counter value. In the case of a confirmed downlink, if the acknowledge is not received, the Application Server is notified and may decide to transmit a new confirmed frame.

Uplink frames

Uplink confirmed and unconfirmed frames are transmitted `NbTrans` times (see Section 5.2) unless a valid Class A downlink is received following one of the transmissions. The `NbTrans` parameter can be used by a Network Server to control the redundancy of end-device uplinks to achieve a given quality of service. An end-device SHALL perform frequency hopping as usual between repeated transmissions. It SHALL wait after each repetition until the receive windows have expired. The delay between retransmissions is at the discretion of the end-device and MAY be different for each end-device. When an end-device has requested an ACK from the Network but has not yet received it, it SHALL wait RETRANSMIT_TIMEOUT seconds after RECEIVE_DELAY2 seconds have elapsed after the end of the previous uplink transmission before sending a new uplink (repetition or new frame). The RETRANSMIT_TIMEOUT delay is not required between unconfirmed uplinks, or after the ACK has been successfully demodulated by the end-device. The RETRANSMIT_TIMEOUT value is given in [RP002]. The retransmission backoff mechanism, defined in Section 7, may also extend the interval between retransmissions, if applicable.

End-devices SHALL stop any further retransmission of an uplink confirmed frame if a corresponding downlink acknowledgment frame is received.

End-devices SHALL also stop any further retransmission of an uplink unconfirmed frame whenever a valid downlink frame is received in a Class A receive window.

If the Network receives more than `NbTrans` transmissions of the same uplink frame having the ADR bit set, this may indicate a replay attack or a malfunctioning end-device, and therefore the Network SHALL silently discard the extra frames.

> **Note:** A Network Server that detects a replay attack may take additional measures, such as reducing the `NbTrans` parameter to 1 or discarding uplink frames received over a channel that was already used by an earlier transmission of the same frame, or by some other unspecified mechanism.

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

The frame pending bit (FPending) is used only in downlink communication. For Class A end-devices, FPending indicates that the Network Server has more data pending to be sent and therefore the end-device MAY send an uplink frame as soon as possible. For Class B end-devices, FPending indicates the priority of which conflicting ping slots the end-device SHALL listen to in case of a collision[6].

An example use of the FPending bit is described in Section 16.4.

4.3.1.5 Frame counter (FCnt)

There are two frame counters for each end-device. FCntUp is incremented by an end-device when a data frame is transmitted to a Network Server (uplink). FCntDown is incremented by a Network Server when a data frame is transmitted to an end-device (downlink). The Network Server tracks the uplink frame counter and generates the downlink counter for each end-device.

Whenever an OTAA end-device successfully processes a Join-Accept frame, the frame counters on the end-device (FCntUp) and the Network side (FCntDown) are reset to 0.

For ABP (activation by personalization) end-devices, the frame counters are initialized to 0 by the manufacturer. ABP end-devices SHALL NOT reset the frame counters during the end-device's lifetime. If the end-device is susceptible to losing power during its lifetime (battery replacement, for example), the frame counters SHALL persist during such an event.

Subsequently, FCntUp is incremented with each uplink and FCntDown is incremented with each downlink. At the receiver side, the corresponding counter is kept in sync with the received value, provided the received value has been incremented compared to the current counter value, and the frame MIC field matches the MIC value computed locally using the appropriate network session key. FCntUp SHALL NOT be incremented in the case of multiple transmissions of a confirmed or unconfirmed frame (see NbTrans parameter). Network Servers SHALL drop the application payload of the retransmitted frames and only forward a single instance to the appropriate Application Server.

A first uplink with FCntUp=0 sent by an ABP or an OTAA end-device after a successful Join procedure SHALL be accepted by a Network Server, provided the MIC field is valid. Analogously, a first downlink with FCntDown=0 sent by a Network Server to an ABP or an OTAA end-device after a successful Join procedure SHALL be accepted by the end-device, provided the MIC field is valid.

Frame counters are 32 bits wide. The FCnt field SHALL correspond to the least-significant 16 bits of the 32-bit frame counter (i.e., FCntUp for data frames sent uplink and FCntDown for data frames sent downlink).

The end-device SHALL NOT reuse the same FCntUp value with the same application or network session keys, except for retransmission.

The end-device SHALL NOT process any retransmission of the same downlink frame. Subsequent retransmissions SHALL be ignored without being processed.

[6] Cf. Section 9.2.

> **Note:** This means that an end-device will only acknowledge receipt of a downlink confirmed frame `NbTrans` times. Similarly, an end-device will only generate `NbTrans` uplinks following receipt of a frame with the `FPending` bit set before incrementing its `FCntUp`.

> **Note:** As the `FCnt` field carries only the least-significant 16 bits of the 32-bit frame counter, the server must infer the 16 most-significant bits of the frame counter by observing the traffic.

4.3.1.6 Frame options (`FOptsLen` in `FCtrl`, `FOpts`)

The frame-options length field (`FOptsLen`) in `FCtrl` denotes the actual length of the frame options field (`FOpts`) included in the frame.

`FOpts` transports MAC commands of a maximum length of 15 octets that are piggybacked onto data frames; see Section 5 for a list of valid MAC commands.

If `FOptsLen=0`, the `FOpts` field SHALL be absent. If `FOptsLen≠0`, i.e. if MAC commands are present in the `FOpts` field, the `FPort` value `0` SHALL NOT be used (`FPort` SHALL either not be present or not equal to `0`).

MAC commands SHALL NOT be present in the payload field and the frame options field simultaneously. Should this occur, the end-device SHALL silently discard the frame.

4.3.1.7 Class B enabled bit (`ClassB` in `FCtrl`, uplink only)

The `ClassB` bit set to 1, in an uplink, signals to the Network Server that the end-device has enabled class B and is now ready to receive scheduled downlink pings. Please refer to the Class B section of the document for the Class B specification.

4.3.2 Port field (`FPort`)

If the frame payload field is not empty, the port field SHALL be present. If present, an `FPort` value of `0` indicates that the `FRMPayload` contains only MAC commands; see Section 5 for a list of valid MAC commands. `FPort` values `1..223` (`0x01..0xDF`) are application-specific. `FPort` value `224` is dedicated to the LoRaWAN MAC layer test protocol. `FPort` values `224..255` (`0xE0..0xFF`) are reserved for use and allocation by the LoRa Alliance [TS008].

The purpose of the `FPort` value `224` is to provide a dedicated `FPort` to run MAC compliance test scenarios over-the-air on final versions of end-devices, without having to rely on specific test versions of end-devices for practical aspects. The test is not supposed to be simultaneous with live operations, but the MAC layer implementation of an end-device SHALL be exactly the one used for the normal application. The test protocol is encrypted using the `AppSKey`. This ensures that the Network cannot enable the end-device's test mode without involving the end-device's owner.

> **Note:** If the test runs on an end-device connected to a live network, the way the test application on the Network side learns the `AppSKey` is beyond the scope of the LoRaWAN specification. If the test runs using OTAA on a dedicated test bench (not a live network), the way the `AppKey` is communicated to the test bench for a secured Join process is also beyond the scope of this specification.

The test protocol running at the application layer is defined in [TS009].

Size (octets)	7..22	0..1	0..*N*
`MACPayload`	`FHDR`	`FPort`	`FRMPayload`

Table 10: `MACPayload` format

N is the number of octets of the application payload and SHALL be equal to or less than $N \leq M - 1 -$ (length of `FHDR` in octets), where *M* is the maximum `MACPayload` length.

The valid ranges of both *N* and *M* are region-specific and defined in the “LoRaWAN Regional Parameters” [RP002] document.

4.3.3 MAC frame payload encryption (`FRMPayload`)

If a data frame carries a payload (`FRMPayload`), it SHALL be encrypted before the message integrity code (MIC) is calculated.

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

Key *K* depends on the `FPort` of the data frame:

`FPort`	***K***
0	`NwkSKey`
1..255	`AppSKey`

Table 11: `FPort` list

The encrypted fields are *pld* = `FRMPayload.`

For each data frame, the algorithm defines a sequence of Blocks A_i for $i = 1..k$, where $k = \mathrm{ceil}(\mathrm{len}(pld) / 16)$:

Size (octets)	1	4	1	4	4	1	1
A_i	`0x01`	4 × `0x00`	`Dir`	`DevAddr`	`FCntUp` or `FCntDown`	`0x00`	*i*

Table 12: A_i format

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

The blocks A_i SHALL be encrypted to obtain a sequence S of blocks S_i as follows:

$$S_i = \text{aes128\_encrypt}(K, A_i) \text{ for } i = 1..k$$
$$S = S_1 \mid S_2 \mid .. \mid S_k .$$

Encryption and decryption of the payload SHALL be calculated as follows:

$$\texttt{FRMPayloadPad} = (pld \mid \text{pad}_{16}) \text{ xor } S$$
$$\texttt{FRMPayload} = \texttt{FRMPayloadPad}[0..\text{len}(pld)\text{-}1] .$$

4.4 Message Integrity Code (MIC)

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

$$msg = \texttt{MHDR} \mid \texttt{FHDR} \mid \texttt{FPort} \mid \texttt{FRMPayload} ,$$

where len(msg) denotes the length of the frame in octets.

The MIC SHALL be calculated as follows [RFC4493]:

$$\texttt{CMAC} = \text{aes128\_cmac}(\texttt{NwkSKey}, B_0 \mid msg)$$
$$\texttt{MIC} = \texttt{CMAC}[0..3] ,$$

where block B_0 is defined as follows:

Size (octets)	1	4	1	4	4	1	1
B_0	`0x49`	4 × `0x00`	`Dir`	`DevAddr`	`FCntUp` or `FCntDown`	`0x00`	len(msg)

Table 13: B_0 format

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 a Network Server and the MAC layer of an end-device. MAC layer commands are never visible to the Application Server, nor to the application running on the end-device.

A single data frame MAY 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 set to `0`. Piggybacked MAC commands SHALL always be sent without encryption and SHALL NOT exceed 15 octets. MAC commands sent as `FRMPayload` SHALL always be encrypted and SHALL NOT exceed the maximum `FRMPayload` length.

> **Note:** MAC commands whose content shall be encrypted must be sent in the `FRMPayload` of a separate data frame.

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 End-device	Transmitted by Network Server	Brief Description
0x02	***LinkCheckReq***	x		Used by an end-device to validate its connectivity to a network.
0x02	***LinkCheckAns***		x	Answers ***LinkCheckReq***. Contains the received signal power estimation, which indicates the quality of reception (link margin) to the end-device.
0x03	***LinkADRReq***		x	Requests the end-device to change data rate, TX power, redundancy, or channel mask.
0x03	***LinkADRAns***	x		Acknowledges ***LinkADRReq***.
0x04	***DutyCycleReq***		x	Sets the maximum aggregated transmit duty cycle of an end-device.
0x04	***DutyCycleAns***	x		Acknowledges ***DutyCycleReq***.
0x05	***RXParamSetupReq***[7]		x	Sets the reception slot parameters.
0x05	***RXParamSetupAns***	x		Acknowledges ***RXParamSetupReq.***
0x06	***DevStatusReq***		x	Requests the status of the end-device.
0x06	***DevStatusAns***	x		Returns the status of the end-device, namely its battery level and its radio status.
0x07	***NewChannelReq***		x	Creates or modifies the definition of a radio channel.
0x07	***NewChannelAns***	x		Acknowledges ***NewChannelReq***.
0x08	***RXTimingSetupReq***[7]		x	Sets the timing of the reception slots.
0x08	***RXTimingSetupAns***	x		Acknowledges ***RXTimingSetupReq***.
0x09	***TXParamSetupReq***[7]		x	Used by a Network Server to set the maximum allowed dwell time and MaxEIRP of end-device, based on local regulations.
0x09	***TXParamSetupAns***	x		Acknowledges ***TXParamSetupReq***.
0x0A	***DlChannelReq***[7]		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	***DlChannelAns***	x		Acknowledges ***DlChannelReq***.
0x0B to 0x0C	RFU			
0x0D	***DeviceTimeReq***	x		Used by an end-device to request the current GPS time.
0x0D	***DeviceTimeAns***		x	Answers ***DeviceTimeReq***.
0x0E to 0x0F	RFU			
0x10 to 0x1F	Class B commands (cf. Sections 12).			
0x20 to 0x2F	Reserved for Class C commands.			
0x30 to 0x7F	RFU			
0x80 to 0xFF	Proprietary	x	x	Reserved for proprietary network command extensions.

Table 14: MAC commands

[7] This command has a different acknowledgment mechanism as described in the command definition.

MAC Commands which require an answer from the Network expire after the Class A receive windows have elapsed.

MAC commands are answered/acknowledged by the receiving end in the same order they were transmitted. The answer to each MAC command is sequentially added to a buffer. All MAC commands received in a single frame SHALL be answered in a single frame, which means that the buffer containing the answers SHALL be sent in a single frame.

If the transmitter has a combination of application payload and MAC answers, or new MAC commands to send and they cannot fit in the same frame, the priority for including information in the frame is shown below. Within a single frame, a transmitter SHALL send all higher-priority information before sending any lower- priority information.

Priority Level	Information type
Highest	MAC answers
	New MAC commands
Lowest	Application payload

Table 15: Transmit data insertion prioritization

> **Note:** MAC answers are defined as MAC commands sent by the transmitter in response to received MAC commands. New MAC commands are defined as MAC commands sent by the transmitter but not in response to a received MAC command.

If the MAC command buffer is too large to fit in the frame, the transmitter SHALL truncate the buffer at the end of the last MAC command that is able to fit within the frame. In all cases, the full list of MAC commands SHALL be executed by the receiver, even if the buffer containing the MAC answers must be truncated.

> **Note**: When receiving a truncated MAC answer, a Network Server may retransmit the MAC commands that could not be answered. The Network Server may elect to send a list of MAC commands, which cannot be answered in a single frame, in order to transition the end-device rapidly to an optimal configuration.

> **Note:** In general, the transmitter will reply once to a MAC command. If the answer is lost, the original sender has to resend the command. The original sender decides that the command must be resent when it receives a new frame that does not contain the answer. Only the ***RXParamSetupReq***, ***RXTimingSetupReq, TXParamSetupReq,*** and ***DlChannelReq*** commands impact the downlink parameters and therefore have a different acknowledgment mechanism as described in their corresponding sections.

> **Note:** When a MAC command is initiated by an end-device, the Network Server may only send the acknowledgment/answer in the RX1/RX2 windows immediately following the request. If the answer is not received

> in that slot, the end-device is free to implement any retry mechanism it requires.

> **Note:** The length of a MAC command is variable and is determined during decoding. 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 define out-of-band the lowest common LoRaWAN version of the Network Server and end-device. Without such knowledge, the order of MAC commands shall be according to the version of the LoRaWAN specification that introduced a MAC command for the first time. This way, all MAC commands up to the implemented version of the LoRaWAN specification can be processed even in the presence of MAC commands specified in newer versions of the LoRaWAN specification.

5.1 Link Check Commands (*LinkCheckReq, LinkCheckAns*)

End-devices and Network Servers SHALL implement these commands.

An end-device MAY use the ***LinkCheckReq*** command to 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, the Network Server SHALL respond with a ***LinkCheckAns*** command.

Size (octets)	1	1
***LinkCheckAns* payload**	`Margin`	`GwCnt`

Table 16: *LinkCheckAns* payload format

The demodulation margin (`Margin`) is an 8-bit unsigned integer in the range of 0..254, which indicates the link margin in dB of the most recently transmitted ***LinkCheckReq*** command. A value of 0 means that the frame was received at the demodulation floor (0 dB or no margin) whereas a value of 20, for example, means that the frame reached the best gateway 20 dB above the demodulation floor. The value 255 is reserved.

The gateway count (`GwCnt`) is the number of gateways that received the most recent ***LinkCheckReq*** command.

5.2 Link ADR Commands (*LinkADRReq, LinkADRAns*)

End-devices and Network Servers SHALL implement these commands.

A Network Server MAY use the ***LinkADRReq*** command to request that an end-device performs a rate adaptation.

Size (octets)	1	2	1
***LinkADRReq* payload**	DataRate_TXPower	ChMask	Redundancy

Table 17: *LinkADRReq* payload format

Bits	[7:4]	[3:0]
DataRate_TXPower	DataRate	TXPower

Table 18: DataRate_TXPower field format

The requested date rate (DataRate) and TX output power (TXPower) are region-specific and encoded as indicated in the "LoRaWAN Regional Parameters" [RP002] document. The TX output power indicated in the command is to be considered the maximum TX power at which the end-device may operate. An end-device SHALL acknowledge the successful receipt of a command that specifies a higher TX power than it is capable of using. In that case, the end-device SHALL operate at its maximum possible power. An end-device will negatively acknowledge a command that specifies a lower TX power than the end-device is capable of using. In that case, the end-device SHALL operate at its previously configured TX power. The value 0xF (decimal 15) of either DataRate or TXPower means that the end-device SHALL ignore that field and keep the current parameter values. An end-device SHALL support a minimum power control range, such that it can operate from its maximum TX power down to the max (2dBm, maximum TX power – 14dB). It is RECOMMENDED that an end-device supports a minimum TX power of +2dBm.

> **Note:** In case of good RF conditions, the Network should slowly lower an end-device's TX power and/or raise the end-device's data rate to avoid making drastic changes that may strand the end-device.

The channel mask (ChMask) SHALL encode the channels usable for uplink access as follows with bit 0 corresponding to the LSB:

Bits	Usable channels
0	Channel 1
1	Channel 2
..	..
15	Channel 16

Table 19: Channel mask format

If a bit in the ChMask field is set to 1, the corresponding channel SHOULD be used for uplink transmissions if this channel allows the data rate currently used by the end-device. A bit set to 0 means that the corresponding channel SHALL NOT be used.

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

Table 20: Redundancy field format

In the Redundancy bits, the NbTrans field is the number of transmissions for each uplink frame. This applies to both confirmed and unconfirmed uplink frames. The default value is 1, which corresponds to a single transmission of each frame. The valid range is [1:15]. If an NbTrans value of 0 is received, the end-device SHALL use the default value. This field MAY be used by a Network Server to control the redundancy of the uplink transmissions (retransmissions) to obtain a given quality of service. The end-device performs frequency hopping for retransmissions as usual, and it waits as usual after each retransmission until RX2 has expired. Whenever a downlink frame is received during either RX1 or RX2, the end-device SHALL stop any further retransmission of that same uplink frame.

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 turn all channels on or off globally using a specific modulation. The meaning of ChMaskCntl is region-specific and defined in the "LoRaWAN Regional Parameters" [RP002] document.

A Network Server MAY include multiple ***LinkADRReq*** commands within a single downlink frame. For the purpose of configuring the end-device channel mask, the end-device SHALL process all contiguous ***LinkADRReq*** commands in the order present in the downlink frame as a single atomic block command. The end-device SHALL accept or reject all Channel Mask controls in the contiguous block and SHALL provide consistent channel mask ACK status indications for each command in the contiguous block in each ***LinkADRAns*** command, reflecting the acceptance or rejection of this atomic channel mask setting. The end-device SHALL only process the DataRate, TXPower and NbTrans from the last ***LinkADRReq*** command in the contiguous block, as these settings govern the end-device global state for these values. The end-device SHALL provide consistent ACK status in each ***LinkADRAns*** command reflecting the acceptance or rejection of these final settings.

The channel frequencies are region-specific and defined [RP002]. An end-device SHALL answer a ***LinkADRReq*** with a ***LinkADRAns*** command.

Size (octets)	1
***LinkADRAns* payload**	Status

Table 21: *LinkADRAns* payload format

Bits	[7:3]	2	1	0
Status	RFU	PowerACK	DataRateACK	ChannelMaskACK

Table 22: Status field format

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

	`Bit=0`	**`Bit=1`**
`ChannelMaskACK`	The channel mask enables a yet undefined channel or the channel mask required all channels to be disabled or the channel mask is incompatible with the resulting data rate or TX power. 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.
`DataRateACK`	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.
`PowerACK`	The end-device is unable to operate at or below the requested power level. The command was discarded and the end-device state was not changed.	The power level was successfully set.

Table 23: *LinkADRAns* `Status` bits signification

If any of those three bits equals `0`, the command did not succeed, and the end-device SHALL keep its previous state.

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

End-devices and Network Servers SHALL implement these commands.

The ***DutyCycleReq*** command can be used by the Network 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 (octets)	1
***DutyCycleReq* payload**	`DutyCyclePL`

Table 24: *DutyCycleReq* payload format

`Bits`	`7:4`	`3:0`
`DutyCyclePL`	`RFU`	`MaxDutyCycle`

Table 25: `DutyCyclePL` field format

The maximum end-device transmit duty cycle allowed is

$$\text{Aggregated duty cycle} = 1/2^{\texttt{MaxDutyCycle}}$$

The valid range for `MaxDutyCycle` is `[0:15]`. A value of `0` corresponds to 100% duty cycle, therefore “no duty cycle limitation” except the one set by the regional regulation.

> **Note:** when applying a ***DutyCycleReq*** command, the end-device will use whichever is the lowest of either the region limitation for the specific sub-band, or the value from the restriction defined by ***DutyCycleReq*** which is aggregated over all sub-bands.

When `MaxDutyCycle` is different than `0`, an end-device SHALL respect a silence period T_{off} before transmitting a frame on any channel. This silence ensures that the duty cycle limit is met even over short observation windows. It is computed as:

$$Toff = TimeOnAir * (2^{MaxDutyCycle} - 1)$$

> **Note:** the ***DutyCycleReq*** command is used by the Network Server for traffic shaping, to limit the transmissions from a given end-device. It is calculated by the above formula, which may be different than the duty cycle measurement methods defined by the regulations which have such a limitation.

An end-device SHALL answer a ***DutyCycleReq*** with a ***DutyCycleAns*** command. The ***DutyCycleAns*** MAC reply contains no payload.

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

End-devices and Network Servers SHALL implement these commands.

The ***RXParamSetupReq*** command allows a change to the frequency and the data rate set for RX2 following each uplink. The command also allows an offset to be programmed between the uplink and the RX1 slot downlink data rates.

Size (octets)	1	3
***RXParamSetupReq* payload**	`DLSettings`	`Frequency`

Table 26: *RXParamSetupReq* payload format

Bits	7	6:4	3:0
DLSettings	RFU	RX1DROffset	RX2DataRate

Table 27: `DLSettings` field format

The RX1 data-rate offset (`RX1DROffset`) field sets the offset between the uplink data rate and the downlink data rate used to communicate with the end-device on RX1. The default offset is `0`. The offset takes into account the maximum power density constraints for gateways in some regions and balances the uplink and downlink radio link margins.

The RX data rate (`RX2DataRate`) field defines the data rate of a downlink using the second receive window following the same convention as the ***LinkADRReq*** command. For example, `0` means DR0/125 kHz. 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 SHALL be used by the end-device to acknowledge the receipt of a ***RXParamSetupReq*** command. The ***RXParamSetupAns*** command SHALL be added in the `FOpts` field (if `FPort` is either missing or >`0`) or in the `FRMPayload` field (if `FPort=0`) of all uplinks until a Class A downlink is received by the end-device. This guarantees that, even in the case of an uplink frame loss, the Network is always aware of the downlink parameters used by the end-device.

Following the transmission of a ***RXParamSetupReq*** command that modifies RX2 (`Frequency` or `RX2DataRate` fields), the Network Server SHALL NOT transmit a Class C downlink before it has received a valid uplink frame containing ***RXParamSetupAns***.

> **Note:** An end-device that expects to receive Class C downlink frames will send an uplink frame as soon as possible after receiving a valid RXParamSetupReq that modifies RX2 (Frequency or RX2DataRate fields).

The payload contains a single `Status` octet.

Size (octets)	1
***RXParamSetupAns* payload**	`Status`

Table 28: *RXParamSetupAns* payload format

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

Bits	7:3	2	1	0
Status	RFU	RX1DROffsetACK	RX2DataRateACK	ChannelACK

Table 29: `Status` field format

	Bit=0	Bit=1
ChannelACK	The frequency requested is not usable by the end-device.	RX2 slot channel was successfully set
RX2DataRateACK	The data rate requested is unknown to the end-device.	RX2 slot data rate was successfully set
RX1DROffsetACK	The uplink/downlink data rate offset for RX1 slot is not within the allowed range	RX1 data-rate offset was successfully set

Table 30: *RX2SetupAns* Status bits signification

If any of the three bits is equal to 0, the command did not succeed, and the end-device SHALL keep its previous state.

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

End-devices and Network Servers SHALL implement these commands.

A Network Server can use the ***DevStatusReq*** command to request status information from an end-device. The command has no payload. If a ***DevStatusReq*** is received by an end-device, it SHALL respond with a ***DevStatusAns*** command.

Size (octets)	1	1
***DevStatusAns* payload**	Battery	RadioStatus

Table 31: *DevStatusAns* payload format

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

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

Table 32: Battery-level decoding

The RadioStatus field contains the signal-to-noise ratio (SNR) information encoded in the six lowest bits in dB rounded to the nearest integer value for the last successfully received ***DevStatusReq*** command. It is a signed integer with a minimum value of −32 and a maximum value of 31.

Bits	7:6	5:0
RadioStatus	RFU	SNR

Table 33: Status field format

5.6 Creation / Modification of a Channel (*NewChannelReq*, *NewChannelAns, DlChannelReq, DlChannelAns*)

End-devices and Network Servers SHALL implement these commands, unless an end-device is operating in a region where a fixed channel plan is defined, in which case these commands SHALL NOT be implemented. Please refer to "LoRaWAN Regional Parameters" [RP002] for applicable regions.

A Network Server can use the ***NewChannelReq*** command either to create a new bidirectional channel or to modify the parameters of an existing one. The command sets the center frequency of the new channel and the range of uplink data rates that are usable on this channel:

Size (octets)	1	3	1
***NewChannelReq* payload**	`ChIndex`	`Frequency`	`DRRange`

Table 34: *NewChannelReq* payload format

The channel index (`ChIndex`) is the index of the channel being created or modified. Depending on the region and frequency band, the "LoRaWAN Regional Parameters" [RP002] document imposes default channels, which SHALL be common to all end-devices and SHALL NOT be modified by the ***NewChannelReq*** command. If the number of default channels is *N*, the default channels go from 0 to *N*−1, and the acceptable range for `ChIndex` is *N* to 15. An end-device SHALL be able to handle at least 16 different channel definitions. In certain regions, the end-device SHALL have to store more than 16 channel definitions.

The `Frequency` field is a 24-bit unsigned integer. The actual channel frequency (in Hz) is 100 × `Frequency`, whereby values representing frequencies below 100 MHz are reserved for future use. This allows the frequency of a channel to be set anywhere from 100 MHz to 1.67 GHz in increments of 100 Hz. A `Frequency` value of `0` disables the channel. The end-device SHALL check that the frequency is allowed by its radio hardware and SHALL NOT set the channel frequency bit in the `Status` field of the answer if the end-device cannot use this frequency (see below).

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

Bits	7:4	3:0
DRRange	MaxDR	MinDR

Table 35: `DRRange` field format

The minimum data rate (`MinDR`) subfield designates the lowest uplink data rate allowed on this channel. The maximum data rate (`MaxDR`) designates the highest uplink data rate. The mapping of data rate index to physical layer is defined in [RP002] for each region.

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

The end-device SHALL acknowledge the receipt of a ***NewChannelReq*** by returning a ***NewChannelAns*** command. The payload of this frame contains the following information:

Size (octets)	1
***NewChannelAns* payload**	Status

Table 36: *NewChannelAns* payload format

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

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

Table 37: `Status` field format

	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 capabilities of the end-device
Channel frequency ok	The end-device cannot use this frequency	The end-device is able to use this frequency.

Table 38: *NewChannelAns* `Status` bits signification

If either of those bits equals `0`, the command did not succeed, and the end-device SHALL keep its previous state.

The ***DlChannelReq*** command allows the Network to associate a different downlink frequency with the RX1 slot. This command is applicable for all regions supporting the ***NewChannelReq*** command (e.g., EU863-870 and CN779-787, but not US902-928 or AU915-928). In regions where that command is not defined, the end-device SHALL silently drop it.

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

Size (octets)	1	3
***DlChannelReq* payload**	ChIndex	Frequency

Table 39: *DlChannelReq* payload format

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

The frequency (`Frequency`) field is a 24-bit unsigned integer. The actual downlink frequency (in Hz) is 100 × `Frequency`, where values representing frequencies below 100 MHz are reserved for future use. The end-device SHALL check that the frequency is allowed by its radio hardware and return an error otherwise.

> **Note:** to revert the RX1 frequency to the default value, a Network Server can send another ***DlChannelReq*** with the same value as uplink frequency.

If the ***DlChannelReq*** command is defined in the region where the end-device is operating, the end-device SHALL acknowledge the receipt of a ***DlChannelReq*** by returning a ***DlChannelAns*** command. The ***DlChannelAns*** command SHALL be added to the `FOpts` field (if `FPort` is either missing or >`0`) or to the `FRMPayload` field (if `FPort=0`) of all uplinks until a Class A downlink is received by the end-device. This guarantees that, even in the case of an uplink frame loss, the Network is always aware of the downlink frequencies used by the end-device.

The payload of this frame contains the following information:

Size (octets)	1
***DlChannelAns* payload**	`Status`

Table 40: *DlChannelAns* payload format

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

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

Table 41: `Status` field format

	`Bit=0`	**`Bit=1`**
`Channel frequency ok`	The end-device cannot use this frequency	The end-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 42: *DlChannelAns* `Status` bits signification

If either of those bits equals `0`, the command did not succeed, and the end-device SHALL keep its previous state.

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

End-devices and Network Servers SHALL implement these commands.

A Network Server can use the ***RXTimingSetupReq*** command to configure the delay between the end of the TX uplink transmission and the opening of RX1. RX2 always opens 1s after RX1 [RP002].

Size (octets)	1
***RXTimingSetupReq* payload**	`RxTimingSettings`

Table 43: *RXTimingSetupReq* payload format

The `RxTimingSettings` field specifies the delay. The field is split in two 4-bit indexes:

Bits	7:4	3:0
RxTimingSettings	RFU	Del

Table 44: **RxTiminingSettings** field format

The delay is expressed in seconds. `Del 0` is mapped to 1s.

Del	**Delay [s]**
0	1
1	1
2	2
3	3
…	…
15	15

Table 45: **Del** mapping

An end-device SHALL answer ***RXTimingSetupReq*** with ***RXTimingSetupAns*** , which has no payload.

The ***RXTimingSetupAns*** command SHALL be added to the `FOpts` field (if `FPort` is either missing or greater than `0`) or to the `FRMPayload` field (if `FPort=0`) of all uplinks until a Class A downlink is received by the end-device. This guarantees that, even in the case of an uplink frame loss, the Network is always aware of the downlink parameters used by the end-device.

5.8 End-Device Transmit Parameters (*TXParamSetupReq*, *TXParamSetupAns*)

A Network Server MAY use the ***TXParamSetupReq*** command to notify the end-device of the maximum allowed dwell time, i.e. the maximum continuous transmit time of a packet over the air, as well as the maximum allowed end-device Effective Isotropic Radiated Power (EIRP).

Size (octets)	1
***TXParamSetup* payload**	`EIRP_DwellTime`

Table 46: *TxParamSetup* payload format

The structure of `EIRP_DwellTime` field is described below:

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

Table 47: `MaxDwellTime` field format

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

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

Table 48: Maximum EIRP encoding

The maximum EIRP corresponds to an upper bound on the end-device's radio TX power. The end-device is not required to transmit at that power but SHALL NOT radiate more than the specified EIRP.

Bits 4 and 5 define the maximum uplink and downlink dwell times, respectively, which are encoded as per the following table:

Coded Value	**Dwell Time**
0	No limit
1	400 ms

Table 49: Maximum dwell time encoding

If this MAC command is implemented (region-specific), the end-device SHALL acknowledge the ***TXParamSetupReq*** command by sending a ***TXParamSetupAns*** command. This ***TXParamSetupAns*** command contains no payload.

> **Note:** when applying a ***TxParamSetupReq*** command, the end-device must use whichever is the lowest of either the region limitation for EIRP or the value encoded in `MaxEIRP`. It must also use whichever is the lowest of either region limitation of dwell time or value encoded in `UplinkDwellTime`.

The ***TXParamSetupAns*** command SHALL be added in the `FOpts` field (if `FPort` is either missing or >`0`) or in the `FRMPayload` field (if `FPort=0`) of all uplinks until a Class A downlink is received by the end-device. This guarantees that, even in the case of an uplink frame loss, the Network is always aware of the downlink parameters used by the end-device.

When this MAC command is used in a region where it is not required, the end-device SHALL NOT process it and SHALL NOT transmit an acknowledgment [RP002].

5.9 End-Device Time Commands (*DeviceTimeReq, DeviceTimeAns*)

End-devices and Network Servers SHALL implement these commands. This MAC command has been introduced in LoRaWAN L2 1.0.3 [TS001-1.0.3].

An end-device MAY use the ***DeviceTimeReq*** command to request the current network time from the Network Server. This allows the end-device to synchronize its internal clock to the Network's clock. This is specifically useful to speed up the acquisition of the Class B beacon. The request has no payload.

A Network Server MAY use the ***DeviceTimeAns*** command to provide the GPS time to the end-device. The time provided is the GPS time at the end of the uplink transmission. The command has a 5-octet payload defined as follows:

Size (octets)	4	1
***DeviceTimeAns* payload**	32-bit unsigned integer: seconds since epoch*	8-bit unsigned integer: fractional-second in $\frac{1}{256}$ s increments

Table 50: *DeviceTimeAns* payload format

The time provided by the Network SHALL have a worst-case accuracy of ±100 ms. The ***DeviceTimeAns*** command SHALL be sent as a Class A downlink (i.e., over RX1/RX2 of the Class A mode).

(*) The GPS epoch (i.e., January 6, 1980 00:00:00 UTC) is used as origin. The `seconds` field is the number of seconds elapsed since the origin. This field increases monotonically by 1 every second. To convert this field to UTC time, the leap seconds SHALL be taken into account.

> **Example:** Friday, 12 February 2016 at 14:24:31 UTC corresponds to 1139322288s since GPS epoch. As of June 2017, the GPS time is 17s ahead of UTC time.

6 End-Device Activation

To participate in a LoRaWAN network, each end-device SHALL be personalized and activated.

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

An end-device SHALL implement either OTAA or ABP, and MAY implement both OTAA and ABP.

6.1 Data Stored in End-Device after Activation

After activation, the following information is stored in the end-device: a device address (`DevAddr`), a network session key (`NwkSKey`), and an application session key (`AppSKey`).

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

The `DevAddr` consists of 32 bits and identifies the end-device within the current network. The `DevAddr` is allocated by the Network Server of the end-device.

Its format SHALL be as follows:

Bits	[31..32−*N*]	[31−*N*..0]
DevAddr	AddrPrefix	NwkAddr

Table 51: DevAddr fields

The variable *N* is an integer within the [`7:25`] range.

The LoRaWAN protocol supports various network address types with different network address space sizes. The variable-size `AddrPrefix` field SHALL be derived from the Network Server's unique identifier `NetID` (see [TS002]) allocated by the LoRa Alliance with the exception of the `AddrPrefix` values reserved for experimental/private networks. The `AddrPrefix` field enables the discovery of the Network Server that had assigned the `DevAddr`. End-devices and Network Servers that do not respect this rule are considered non-compliant.

The least significant (32-*N*) bits are the network address (`NwkAddr`) of the end-device. They SHALL be arbitrarily assigned by the Network Server.

The following `AddrPrefix` values may be used by private/experimental networks and will not be allocated by the LoRa Alliance.

Private/experimental network reserved `AddrPrefix`
***N* = 7**
`AddrPrefix = 7'b0000000 or AddrPrefix = 7'b0000001`
`NwkAddr` = 25-bit range freely allocated by a Network Server

Table 52: `AddrPrefix` values available for use by private/experimental networks

Please refer to [TS002] for the exact construction of the `AddrPrefix` field and the definition of the various address classes.

6.1.2 Network session key (`NwkSKey`)

`NwkSKey` is a network session key specific to the end-device. It is used by both the Network Server and the end-device to calculate and verify the MIC (message integrity code) of all data frames to ensure data integrity. It is further used to encrypt and decrypt the payload field of MAC-only data frames, where `FPort=0`.

`NwkSKey` SHOULD be stored such that extraction and re-use by malicious actors is prevented.

6.1.3 Application session key (`AppSKey`)

`AppSKey` is an application session key specific to the end-device. It is used by both the Application Server and the end-device to encrypt and decrypt the payload field of application-specific data frames. Application payloads SHALL be encrypted end-to-end between the end-device and the Application Server, but they are integrity-protected only over-the-air and not end-to-end. This means that a Network Server may be able to alter the encrypted content of the data frames in transit (yet without being able to read the plain content). Network servers are considered to be trusted, but it is RECOMMENDED that applications wishing to implement end-to-end confidentiality and integrity protection use additional end-to-end security solutions, which are beyond the scope of this specification.

`AppSKey` SHOULD be stored such that extraction and re-use by malicious actors is prevented.

6.2 Over-the-Air Activation

For over-the-air activation, end-devices SHALL follow a Join Procedure prior to participating in data exchanges with a Network Server. An end-device SHALL initiate a new Join Procedure every time it loses the session context information.

An end-device SHALL be personalized with the following information before it starts the Join procedure: a globally unique end-device identifier (`DevEUI`), the Join Server identifier (`JoinEUI`), and an AES-128 key (`AppKey`).

The `JoinEUI` is described below in Section 6.2.2.

> **Note:** For over-the-air-activation, end-devices are not personalized with any kind of network key. Instead, whenever an end-device joins a network, a network session key specific to that end-device is derived to encrypt and verify transmissions at the network level. This facilitates roaming of end-devices between networks of different providers. Furthermore, using both a network session key and an application session key allows federated Network Servers in which application data cannot be read by the network provider.

6.2.1 End-device identifier (`DevEUI`)

`DevEUI` is a global end-device ID in the IEEE EUI64 address space that uniquely identifies the end-device across roaming networks.

All end-devices SHALL have an assigned `DevEUI`, regardless of which activation procedure is used (i.e., ABP or OTAA).

For OTAA end-devices, `DevEUI` SHALL be stored in the end-device before the Join procedure is executed. For ABP end-devices, `DevEUI` SHOULD be stored in the end-device itself.

> **Note:** It is a recommended practice that `DevEUI` should also be available on an end-device label for the purpose of end-device administration.

6.2.2 Join-Server identifier (`JoinEUI`)

`JoinEUI` is a global application ID in the IEEE EUI64 address space that uniquely identifies the Join-Server that is able to assist in the processing of the Join procedure and the derivation of session keys.

For OTAA end-devices, `JoinEUI` SHALL be stored in the end-device before the Join procedure is executed. `JoinEUI` is not required for ABP-only end-devices.

6.2.3 Application key (`AppKey`)

The `AppKey` is an AES-128 root key specific to the end-device.[8] Whenever an end-device joins a network via over-the-air activation, the `AppKey` is used to derive the session keys `NwkSKey` and `AppSKey` specific to that end-device to encrypt and verify network communication and application data.

An `AppKey` SHALL be stored on an end-device intending to use the OTAA procedure.

An `Appkey` is NOT REQUIRED for ABP-only end-devices.

6.2.4 Join procedure

From an end-device's point of view, the Join procedure consists of two MAC frames exchanged with the server, namely a Join-Request and a Join-Accept.

[8] As all end-devices end up with unrelated application keys specific to each end-device, extracting the `AppKey` from an end-device compromises only that one end-device.

6.2.5 Join-Request frame

The Join procedure is always initiated by the end-device sending a Join-Request frame.

Size (octets)	8	8	2
Join-Request payload	`JoinEUI`	`DevEUI`	`DevNonce`

Table 53: Join-Request payload format

The Join-Request frame contains the `JoinEUI` and `DevEUI` of the end-device followed by a nonce of 2 octets (`DevNonce`).

`DevNonce` is a counter starting at `0` when the end-device is initially powered up and incremented with every Join-Request. A `DevNonce` value SHALL never be reused for a given `JoinEUI` value. If the end-device can be power-cycled, then `DevNonce` SHALL be persistent (e.g., stored in a non-volatile memory). Resetting `DevNonce` without changing `JoinEUI` will cause the Join Server to discard the Join-Requests of the end-device. For each end-device, the Join Server keeps track of the last `DevNonce` value used by the end-device and ignores Join-Requests if `DevNonce` is not incremented.

The message integrity code (`MIC`) value (see Section 4 for MAC frame description) for a Join-Request frame SHALL be calculated as follows:[9]

`CMAC` = aes128_cmac(`AppKey`, `MHDR` | `JoinEUI` | `DevEUI` | `DevNonce`)
MIC = `CMAC`[`0`..`3`]

The Join-Request frame is not encrypted.

The Join-Request frame MAY be transmitted using any data rate and following a random frequency-hopping sequence across the specified Join channels. Join-Request transmission PHY requirements can be region-specific and are detailed in [RP002]. If Join-Request transmission PHY parameters are not available in [RP002], end-devices SHALL transmit Join-Requests across all available channels at the lowest required data rate for the region. Additionally, end-devices SHOULD transmit Join-Requests across all available channels at all required data rates for that region. The intervals between transmissions of Join-Requests SHALL respect the conditions described in Section 7. For each transmission of a Join-Request, the end-device SHALL increment the `DevNonce` value.

6.2.6 Join-Accept frame

The Network SHALL respond to a Join-Request frame with a Join-Accept frame if the end-device is permitted to join the network. The Join-Accept frame is sent like a normal downlink but uses delays JOIN_ACCEPT_DELAY1 or JOIN_ACCEPT_DELAY2 (instead of RECEIVE_DELAY1 and RECEIVE_DELAY2, respectively). The channel frequency and data rate used for these two receive windows are identical to the ones used for the RX1 and RX2 receive windows described in Section 3.3 and SHALL use the default values defined in the Regional Parameters [RP002] specification.

A Join-Accept response SHALL NOT be given to the end-device if the Join-Request is not accepted.

[9] [RFC4493]

The Join-Accept frame contains a Join-Server nonce (JoinNonce) of 3 octets, a network identifier (NetID), an end-device address (DevAddr), downlink configuration settings (DLSettings), a delay between TX and RX (RXDelay) and an OPTIONAL list of network parameters (CFList) for the Network the end-device is joining. The OPTIONAL CFList is region-specific and is defined in [RP002].

Size (octets)	3	3	4	1	1	(16) optional
Join-Accept payload	JoinNonce	NetID	DevAddr	DLSettings	RXDelay	CFList

Table 54: Join-Accept payload format

JoinNonce is a non-repeating value provided by the Join Server and used by the end-device to derive the two session keys NwkSKey and AppSKey, which SHALL be calculated as follows:[10]

NwkSKey = aes128_encrypt(AppKey, 0x01 | JoinNonce | NetID | DevNonce | pad_{16})
AppSKey = aes128_encrypt(AppKey, 0x02 | JoinNonce | NetID | DevNonce | pad_{16})

The MIC value for a Join-Accept frame SHALL be calculated as follows:[11]

CMAC = aes128_cmac(AppKey, MHDR | JoinNonce | NetID | DevAddr | DLSettings | RXDelay | CFList)
MIC = CMAC[0..3]

The Join-Accept frame itself SHALL be encrypted with the AppKey as follows:

aes128_decrypt(AppKey, JoinNonce | NetID | DevAddr | DLSettings | RXDelay | CFList | MIC)

> **Note:** [TR001] proposes additional behavior for the JoinNonce value in the Join Server to prevent synchronization issues related to the LoRaWAN 1.0.x Join Procedure. Some of the remedies include additional behavior both at the end-device and the Join Server, which are expected to be configured synchronously. See [TR001] for details.

> **Note:** An AES decrypt operation in ECB mode encrypts the Join-Accept frame so that the end-device can use an AES encrypt operation to decrypt the frame. This way, an end-device has to implement only AES encrypt but not AES decrypt.

> **Note:** Establishing these two session keys allows for a federated network infrastructure in which network operators are not able to eavesdrop on application data. The application provider commits to the network operator that it will take the charges for any traffic incurred by

[10] The pad_{16} function appends all-zero octets so that the length of the data is a multiple of 16.
[11] [RFC4493].

> the end-device and retains full control over the `AppSKey` used for protecting its application data.

The format of the `NetID` is described in [TS002].

The `DLSettings` field SHALL contain the downlink configuration:

Bits	7	6:4	3:0
DLSettings	RFU	RX1DROffset	RX2DataRate

Table 55: `DLSettings` field format

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 receive window (RX1). The default offset is `0`. The offset accommodates the maximum power density constraints for gateways in some regions and balances the uplink and downlink radio link margins.

The `RX2DataRate` field sets the downlink data rate that serves to communicate with the end-device on the second receive window (RX2).

The `RXDelay` field sets the downlink RX1 delay and follows the same convention as the `Del` field in the ***RXTimingSetupReq*** command.

Default RX data rates, default RX receive delays and the actual relationship between the uplink and downlink data rate are region-specific and detailed in the "LoRaWAN Regional Parameters" [RP002] document.

The `CFList` parameters, when present, SHALL be used in combination with implicit parameters (defined in [RP002]) to emulate the receipt of existing MAC commands, such as ***NewChannelReq*** or ***LinkADRReq***. The end-device behavior when processing `CFList` SHALL be identical to what would result from processing those MAC commands if they were received in a single downlink frame, after the processing of the non-optional Join-Accept parameters, with the exception that `CFList` processing does not generate a MAC answer. The standard behavior for processing MAC commands is defined in a subsequent section, and potentially adapted by [RP002].

If the Join-Accept frame is received following the transmission of a Join-Request, the end-device SHALL revert to its default channel and RF parameters definitions. All MAC layer parameters (except `RXDelay`, `RX2DataRate`, and `RX1DROffset` that are transported by the Join-Accept frame) SHALL be reset to their default values. If the `CFlist` is present, it is then applied as defined in [RP002].

It is RECOMMENDED that the first uplink that follows the reception of the Join-Accept frame uses the data rate of the successful Join-Request, or a lower data rate. The end-device SHALL then apply the adaptive date-rate control as defined in section 4.3.1.1.

6.2.7 Join procedure completion for Class C

The end-device that expects to receive Class C downlink frames SHALL send a confirmed uplink frame or a frame that requires an acknowledgment as soon as possible after receiving a valid Join-Accept frame. The end-device SHALL continue to send such frames until it receives the first downlink from the Network (while respecting duty cycles, if applicable, and retransmission timers).

The Network Server SHALL NOT transmit a downlink before it has received a first uplink frame.

6.3 Activation by Personalization

Activation by personalization ties an end-device directly to a specific network, thus bypassing the Join-Request – Join-Accept procedure.

Activating an end-device by personalization means that the `DevAddr` and the two session keys `NwkSKey` and `AppSKey` are stored directly in the end-device instead of being derived from `DevEUI`, `JoinEUI` and the `AppKey`. The end-device is equipped with the required information for participating in a specific LoRaWAN network as soon as it is started.

Each end-device SHALL have a unique set of `NwkSKey` and `AppSKey` values. Compromising the keys of one end-device SHALL NOT compromise the security of the communications of other end-devices. The process to build those keys SHALL be such that the keys cannot be derived in any way from publicly available information such as the end-device address or `DevEUI`.

Upon first boot and following a reset, personalized end-devices SHALL have all available channels for that region enabled and SHOULD use all required data rates for that region. Configurations of the end-device by the Network that controls downlink connectivity (controlled by ***RXParamsSetupReq***, ***DlChannelReq***, ***RXTimingSetupReq***, and ***TXParamSetupReq***) SHALL be persisted by the end-device, even after a reset.

Frame counter values SHALL be used only once in all invocations of a same key with the CCM* (Counter with CBC Message Authentication Code) mode of operation [IEEE802154]. Therefore, re-initialization of an ABP end-device frame counters is forbidden. ABP end-devices SHALL store the frame counters persistently (e.g., in non-volatile memory).

> **Note:** ABP end-devices use the same session keys throughout their lifetime (i.e., no rekeying is possible). Therefore, it is recommended that OTAA end-devices be used for higher security applications.

7 Retransmissions Backoff

Uplink frames that

- require an **acknowledgment or answer** from the Network or an Application Server and are **retransmitted** by the end-device if the acknowledgment or answer is not received.

and

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

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

> **Note:** A typical example of such an uplink frame is a Join-Request if the implementation of a group of end-devices decides to reset the MAC layer in the case of a network outage. The entire group of end-devices will start broadcasting Join-Request uplinks and will stop only upon receiving a Join-Accept 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 end-device (for example using a pseudo-random generator seeded with the end-device's address). The transmission duty-cycle of such a frame SHALL respect local regulations 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 < 36 s per hour	1% duty cycle
Aggregated during the next 10 hours	$T_{0+1} < t < T_{0+11}$	Transmit time < 36 s per 10 h	0.1% duty cycle
After the first 11 hours, aggregated over 24 h, where N refers to days starting at 0	$T_{0+11} + N \times$ (24 hours/day) $< t < T_0 + 35 + N \times$ (24 hours/day), $N \geq 0$	Transmit time < 8.7 s per 24 h	0.01% duty cycle

Table 56: Transmit duty-cycle limitations

1602
CLASS B – BEACON

8 Introduction to Class B

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

End-devices SHOULD implement Class B operation when there is a requirement to open receive windows at fixed time intervals for the purpose of enabling network-initiated downlink frames. Class B-capable end-devices SHALL NOT enable Class B and Class C operation concurrently.

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

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

> **Note:** The decision to switch from Class A to Class B comes from the application layer of the end-device. If this switch from Class A to Class B has to be controlled from the network side, the customer application must use one of the end-device's Class A uplinks to send back a downlink to the application layer, and it needs the application layer on the end-device to recognize this request—this process is not managed at the LoRaWAN level.

8.1 Principle of Synchronous Network-initiated Class B Downlinks

For a network to support end-devices of Class B, the Network SHALL broadcast a beacon that provides a timing reference to end-devices. Based on this timing reference, the Class B-enabled end-devices SHALL periodically open receive windows, hereafter called ping slots, which can be used by the Network 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. For this reason, if an end-device moves and detects a change in the identity advertised in the received beacon, it SHALL send an uplink to the Network Server so that the server can update the downlink routing path database.

When enabling Class B mode, an end-device SHALL use the defined values for the following parameters:

- Default ping-slot periodicity
- Default ping-slot data rate
- Default ping-slot channel.

These parameters have default values defined in the "Regional Parameters Specification" [RP002] and MAY be updated via Class B MAC commands (cf. Section 12).

All end-devices start and join the network as Class A end-devices with Class B disabled. The end-device application can then decide to enable Class B. Class B-capable end-devices still implement all functionalities of Class A end-devices. In particular, Class B-enabled end-devices SHALL respect the Class A RX1 and RX2 receive window definition following every uplink (cf. Section 3.3).

Class B is enabled by the following process:

- The end-device application requests the LoRaWAN layer to enable Class B mode. The LoRaWAN layer in the end-device searches for a beacon. To accelerate beacon discovery, the LoRaWAN layer MAY use the ***DeviceTimeReq*** MAC command.
- Once the end-device has found a beacon, it MAY enable Class B mode.
- Once Class B is enabled, the MAC layer SHALL set to 1 the `ClassB` bit of the `FCtrl` field of every uplink frame transmitted to remain Class B-enabled. This bit signals to the server that the end-device has enabled Class B.
- The MAC layer SHOULD autonomously schedule a reception slot for each beacon and each ping slot. The end-device SHALL take 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 SHALL be processed similarly to a downlink as described in the LoRaWAN Class A specification.
- A mobile end-device SHOULD 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 SHALL appropriately set the `ClassB` bit to `1` in the frame's `FCtrl` field. This can be done more efficiently if the end-device detects that it is moving by analyzing the beacon content. In that case, to avoid systematic uplink collisions, the end-device SHALL apply a random delay (as defined in Section 13.6) between having received the beacon and transmitting the uplink.
- During any Class A downlink, the Network Server MAY change the end-device's ping-slot downlink frequency or data rate by sending a ***PingSlotChannelReq*** MAC command and receiving the corresponding ***PingSlotChannelAns***.
- The end-device MAY change the periodicity of its ping slots at any time. To do so, it SHALL temporarily disable Class B operation (unset `ClassB` bit in its uplink frames) and send a ***PingSlotInfoReq*** to the Network Server. Once this command is acknowledged, the end-device MAY re-enable Class B operation with the new ping-slot periodicity.
- If no beacon has been received for a given period (as defined in Section 12.2), synchronization with the Network is lost. The end-device LoRaWAN layer SHALL stop setting the `ClassB` bit in all uplinks, which informs the Network Server that the end-device has disabled Class B mode.

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

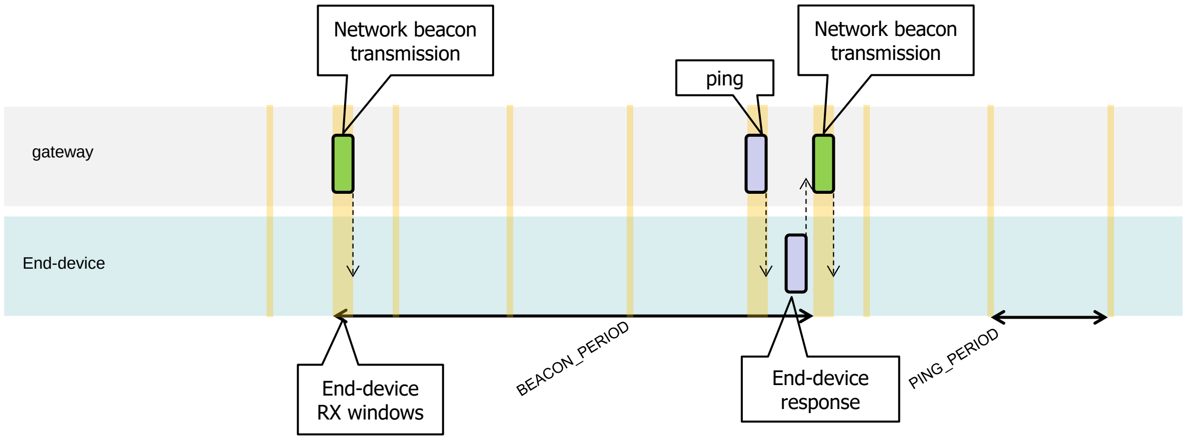

Figure 3: Example of beacon reception slot and ping slots

In this example, given a beacon period of 128s, the end-device also opens a ping-slot reception window every 32s. Most of the time, this ping slot is not used by the Network Server and therefore the end-device reception window is closed as soon as the radio transceiver has assessed that no preamble is present on the radio channel. If a preamble is detected, the radio transceiver will stay on until the downlink frame is demodulated. The MAC layer will then process the frame, check that its address field matches the end-device address and that the MIC is valid before forwarding it to the application layer. The end-device response shown in this example is optional, depending on the downlink, and, if present, this is a Class A uplink.

9 Class B Frame Formats

9.1 Uplink Frames

The uplink frames in Class B mode are the same as the Class A uplinks. The `ClassB` bit SHALL be set to `1` in an uplink to signal the Network Server that the end-device is Class B-enabled and is ready to receive scheduled downlink pings.

9.2 Downlink Frames

The `FPending` bit for Class B downlink frames signals that, in the event of a ping-slot collision between multiple Class B ping slots, the ping-slot sequence whose `FPending` bit was set first will take priority. If the `FPending` bit has been set for multiple Class B ping-slot sequences, they will take priority over ping-slot sequences for which `FPending` has not been set. Further prioritization can be determined based on whether the downlink ping slot is used for multicast or unicast frames.

Priority	**Type of Class B Downlink**
Highest	Multicast with `FPending` previously set
	Unicast with `FPending` previously set
	Multicast with `FPending` not previously set
Lowest	Unicast with `FPending` not previously set

Table 57: `FPending` Class B prioritization

If there are multiple colliding ping-slot sequences at the same priority level as shown in Table 57, the highest unicast or multicast `DevAddr` SHALL take priority.

9.3 Downlink Ping Frames

A downlink ping frame uses the same format as a Class A downlink frame but might follow a different channel frequency or data rate plan.

Frames can be unicast or multicast. Unicast frames are sent to a single end-device, whereas multicast frames are sent to multiple end-devices. All end-devices of a multicast group SHALL share the same multicast address and associated encryption keys. A Class B-capable end-device SHALL support at least one multicast group and an end-device SHALL NOT transmit using a multicast address.

More precisely, inside a given end-device, a multicast group is defined by the following parameters called the multicast group context:

1. A 4-octet network address of the multicast group, common to all end-devices of the group.
2. The multicast group sessions keys (different for every multicast group, but all end-devices of a given multicast group have the same session keys).
3. The multicast group downlink frame counter.

The LoRaWAN Class B specification does not specify means to set up such a multicast group remotely or to distribute the required multicast key material securely. This can be performed either during the end-device personalization or through the application layer.

> **Example:** The document [RPD1] describes a possible application layer mechanism for over-the-air multicast key distribution.

9.3.1 Unicast downlink ping frame format

The MAC payload of a unicast downlink ping uses the format defined in the Class A specification. The same frame counter is incremented, whether the downlink uses a Class B ping slot or a Class A downlink slot. A downlink ping is processed by the end-device in the same way as a Class A downlink, except for MAC commands and confirmed frames.

A downlink ping SHALL NOT transport any MAC command. If an end-device receives a downlink ping containing a MAC command, either in the `FOpts` field (if `FPort` is either missing or >`0`) or in the `FRMPayload` field (if `FPort=0`), it SHALL silently discard the entire frame.

In case a confirmed downlink ping frame is received, the end-device SHALL NOT answer later than a time period equal to CLASS_B_RESP_TIMEOUT * `NbTrans` + RECEIVE_DELAY2 * (`NbTrans`-1) when the end-device sets its uplink ADR bit, and CLASS_B_RESP_TIMEOUT when the end-device unsets its uplink ADR bit. The default value of CLASS_B_RESP_TIMEOUT is 8s. It can be modified in [RP002], it SHALL not be smaller than RETRANSMIT_TIMEOUT plus the maximum time on air of the uplink frame.

> **Note:** The purpose of this timeout is for the Network Server to know how long to wait for a response from the end-device, before it transmits another confirmed downlink. This ensures that the Network Server can process responses without any ambiguity: there may be only a single confirmed downlink ping pending an acknowledgement.

The uplink frame sent in response MAY be sent up to `NbTrans` times, but no retransmission SHALL occur after the timeout period. It is RECOMMENDED that the end-device transmits each copy of the uplink frame at a random time within CLASS_B_RESP_TIMEOUT.

As this adds a timing requirement compared to responding to Class A downlinks, the end-device may not send an uplink frame within the timeout, for instance if it has a duty cycle limitation. When such an uplink frame is not sent, the end-device SHALL act as if the uplink frame with the ACK bit set was sent.

After sending a confirmed downlink frame sent over pingslot, the network server SHALL NOT send any other confirmed downlink to the end-device until this timeout expires or an uplink frame is received from that end-device.

After sending a Class A confirmed downlink, a network server SHALL NOT send any other confirmed downlink to the end-device until an uplink frame is received from that end-device.

> **Note:** unconfirmed downlink frames sent over ping slots may be sent without a minimum delay between them.

9.3.2 Multicast downlink ping frame format

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

- They SHALL NOT carry MAC commands in the `FOpts` field nor in the payload on port 0 because a multicast downlink does not have the same authentication robustness as a unicast frame. The end-device SHALL discard any multicast frame carrying MAC commands.
- The `ACK` bit SHALL be `0` and the `FType` field SHALL have the value `Unconfirmed Data Down`, see Table 4: MAC frame types. The end-device SHALL discard the multicast frame otherwise.

10 Class B Beacon Acquisition and Tracking

Before enabling Class B operation, the end-device SHOULD first synchronize with the network beacons to align its internal timing reference with the network.

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

A Class B-enabled end-device may be temporarily unable to receive beacons (out of range from the network gateways, presence of interference, etc.). In this event, the end-device SHOULD gradually widen its beacon and ping-slot reception windows to accommodate a possible drift of its internal clock.

> **Note:** For example, an end-device whose internal clock is defined with a precision of ±10 ppm may drift by ±1.3 ms every beacon period.

10.1 Minimal Beaconless Operation Time

In the event of beacon loss, an end-device SHALL be capable of maintaining Class B operation for 2 hours (120 minutes) after it received the last beacon. This temporary Class B operation without beacon is called beaconless operation. It relies on the end-device's own clock to keep time.

During beaconless operation, Class B unicast, multicast and beacon reception slots SHALL be progressively expanded to accommodate the end-device's possible clock drift.

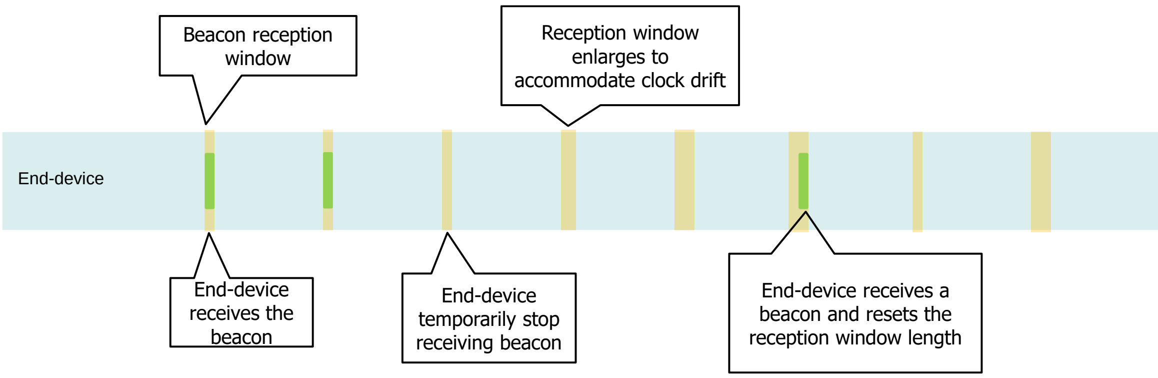

Figure 4: Beaconless temporary operation

10.2 Extension of Beaconless Operation upon Receipt

During the 120-minute time interval described above, the receipt of any beacon directed to the end-device SHOULD extend Class B beaconless operation by another 120 minutes allowing the end-device to correct any timing drift and reset the duration of the receive slots.

> **Note:** An end-device can also use Class B ping-slot downlinks to resynchronize its internal clock.

10.3 Minimizing Timing Drift

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

11 Class B Downlink Slot Timing

11.1 Definitions

To operate successfully in Class B, end-devices SHALL open reception slots at precise instants relative to the infrastructure beacon. This section defines the required timing.

The interval between the start of two successive beacons is called the beacon period. The beacon frame transmission is aligned with the beginning of the BEACON_RESERVED interval. Each beacon is preceded by a BEACON_GUARD time interval, where no ping slot can be placed. The length of the BEACON_GUARD time interval corresponds to the time on air of the longest allowed frame. This is to ensure that a Class B downlink initiated during a ping slot just before the BEACON_GUARD time interval will always have time to finish without colliding with the beacon transmission. The usable time interval for a ping slot therefore spans from the end of the BEACON_RESERVED time interval to the beginning of the next BEACON_GUARD time interval.

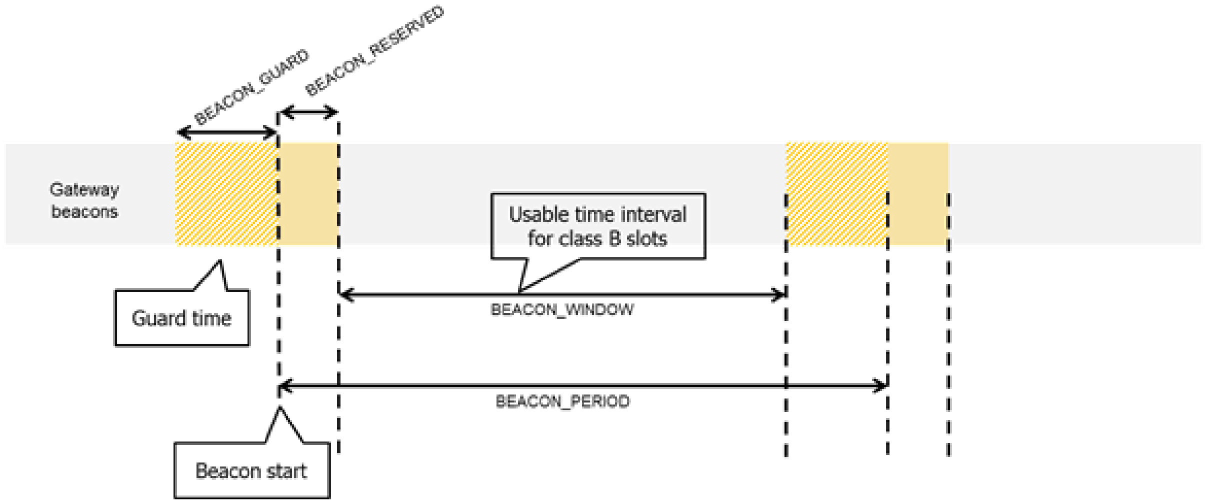

Figure 5: Beacon timing

BEACON_PERIOD	128 s
BEACON_RESERVED	2.120 s
BEACON_GUARD	3.000 s
BEACON_WINDOW	122.880 s

Table 58: Beacon timing

The beacon frame time on air is actually much shorter than the BEACON_RESERVED time interval to allow network management broadcast frames to be appended in the future.

The BEACON_WINDOW time interval is divided into $2^{12} = 4096$ ping slots of 30 ms each, numbered from 0 to 4095.

An end-device using the slot number *N* SHALL turn its receiver on T_{on} s after the start of the beacon, where

$$T_{on} = \text{BEACON\_RESERVED} + N \times 30 \text{ ms}.$$

N is called the slot index.

The latest ping slot starts at BEACON_RESERVED + 4095 × 30 ms = 124,970s after the beacon starts or 3030 ms before the beginning of the next beacon.

11.2 Slot Randomization

To avoid systematic collisions or overhearing problems, the slot index is randomized and changed at every beacon period.

The following parameters are used:[12]

`DevAddr`	Device 32-bit network unicast or multicast address
`PingNb`	Number of ping slots per beacon period. This is a power of 2 integer: $\texttt{PingNb} = 2^{7-\text{Periodicity}}$
`PingPeriod`	Period of the end-device receiver wake-up expressed in number of slots: $\texttt{PingPeriod} = 2^{5+\text{Periodicity}}$
`PingOffset`	Randomized offset computed at each BEACON_PERIOD start. Values can range from 0 to (`PingPeriod-1`)
`BeaconTime`	The time carried in the field `BCNPayload`.
`SlotLen`	Length of a unit ping slot = 30 ms

Table 59: Class B slot randomization algorithm parameters

At each beacon period, the end-device and the server SHALL compute a new pseudo-random offset to align the reception slots. An AES encryption with a fixed key of 16 all-zero octets SHALL be used to randomize:

Key = 16 × `0x00`
Rand = aes128_encrypt(Key, `BeaconTime` | `DevAddr` | pad16)
`PingOffset` = (Rand[0] + Rand[1] × 256) modulo `PingPeriod`

The slots used for this beacon period SHALL be
`PingOffset` + *N* × `PingPeriod` with *N*=[`0:PingNb-1`].

[12] Periodicity is defined in Section 12.1.

The end-device therefore opens receive slots starting at:

First slot	BEACON_RESERVED + `PingOffset` × `SlotLen`
Slot 2	BEACON_RESERVED + (`PingOffset` + `PingPeriod`) × `SlotLen`
Slot 3	BEACON_RESERVED + (`PingOffset` + 2 × `PingPeriod`) × `SlotLen`
…	…

Table 60: Receive-slot starting times

If the end-device simultaneously serves a unicast and one or more multicast slots, this computation SHALL be performed multiple times at the beginning of a new beacon period: once for the unicast address (end-device network address) and once for each multicast group address.

If a Class A RX1 or RX2 receive slot collides with a Class B multicast or unicast slot, the end-device SHALL listen to the Class A RX slot that has priority.

> **Note:** As defined in the Class A section, the end-device does not open the RX2 receive window if a valid unicast downlink addressed to the end-device is received in the RX1 receive window.

The randomization scheme prevents a systematic collision between unicast and multicast slots. If a collision occurs during a beacon period, one is unlikely to occur again during the next beacon period.

12 Class B MAC Commands

All commands described in the Class A specification SHALL be implemented in Class B-capable end-devices. End-devices implementing the Class B specification SHALL further implement the following MAC commands (cf. Table 14: MAC commands).

CID	Command	Transmitted by		Brief Description
		End-device	Network Server	
`0x10`	***PingSlotInfoReq***	x		Used by the end-device to communicate the unicast ping-slot periodicity to the Network Server
`0x10`	***PingSlotInfoAns***		x	Used by the Network to acknowledge a ***PingInfoSlotReq*** command
`0x11`	***PingSlotChannelReq***[13]		x	Used by the Network Server to set the unicast ping channel frequency and data rate of an end-device
`0x11`	***PingSlotChannelAns***	x		Used by the end-device to acknowledge a ***PingSlotChannelReq*** command
`0x12`	***BeaconTimingReq***	x		Deprecated
`0x12`	***BeaconTimingAns***		x	Deprecated
`0x13`	***BeaconFreqReq***		x	Command used by the Network Server to modify the frequency at which the end-device expects to receive a beacon broadcast
`0x13`	***BeaconFreqAns***	x		Used by the end-device to acknowledge a ***BeaconFreqReq*** command

Table 61: Class B MAC command table

MAC Commands which require an answer from the Network expire after the Class A receive windows have elapsed.

12.1 *PingSlotInfoReq*

An end-device MAY use the ***PingSlotInfoReq*** command to inform the server of its unicast ping-slot periodicity. This command SHALL be used only to inform the server of the periodicity of a unicast ping slot. A multicast slot is entirely defined by the application and SHALL NOT use this command.

Size (octets)	1
***PingSlotInfoReq* payload**	`PingSlotParam`

Table 62: *PingSlotInfoReq* payload format

[13] This command has a different acknowledgment mechanism as described in the command definition.

Bits	7:3	[2:0]
PingSlotParam	RFU	Periodicity

Table 63: `PingSlotParam` field format

The `Periodicity` subfield is an unsigned 3-bit integer encoding the ping-slot period currently used by the end-device using the following equations:

$\texttt{pingNb} = 2^{7\text{-}\texttt{Periodicity}}$ and $\texttt{pingPeriod} = 2^{5+\texttt{Periodicity}}$ slots.

The actual ping-slot periodicity will be $0.96 \times 2^{\texttt{Periodicity}}$ s.

- `Periodicity=0` means that the end-device opens a ping slot approximately every 1s during the BEACON_WINDOW interval.
- `Periodicity=7` means that the end-device opens a ping slot approximately every 128s, which is the maximum ping-slot period supported by the LoRaWAN Class B specification.

To change its ping-slot periodicity, an end-device SHALL first revert to Class A. Next it SHALL send the new periodicity through a ***PingSlotInfoReq*** command. Then it SHALL receive an acknowledge from the server through a ***PingSlotInfoAns***. Only then MAY the end-device switch back to Class B with the new periodicity.

This command MAY be concatenated with any other MAC command in the `FOpts` field of `FHDR`, as described in the Class A specification frame format.

Upon receiving this ***PingSlotInfoReq*** command, the Network Server SHALL answer with a ***PingSlotInfoAns*** frame. The MAC payload of this frame is empty.

12.2 *BeaconFreqReq*

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

Octets	3
***BeaconFreqReq* payload**	Frequency

Table 64: *BeaconFreqReq* payload format

The `Frequency` coding is identical to the ***NewChannelReq*** MAC command defined for Class A.

A valid non-zero `Frequency` SHALL force the end-device to listen to the beacon on a fixed frequency channel, even if the default behavior specifies a frequency hopping beacon (i.e. US ISM band).

A value of `0` instructs the end-device that it SHALL use the default beacon frequency plan as defined in Section 13.1. Where applicable, the end-device SHALL resume a frequency hopping beacon search.

Upon receiving this command, the end-device SHALL answer with a ***BeaconFreqAns*** frame. The MAC payload of this frame contains the following information:

Size (octets)	1
***BeaconFreqAns* payload**	`Status`

Table 65: *BeaconFreqAns* payload format

The bits of the `Status` octet have the following meaning:

Bits	7:1	0
Status	RFU	Beacon frequency ok

Table 66: `Status` field format

	Bit=0	**Bit=1**
Beacon frequency ok	The end-device cannot use this frequency, the previous beacon frequency is kept	The beacon frequency has been changed

Table 67: Meaning of beacon frequency bits

12.3 *PingSlotChannelReq*

The server MAY send this command to the end-device to modify the frequency and/or the data rate at which the end-device expects the downlink pings.

Once the Network Server has sent the ***PingSlotChannelReq*** command, it SHALL NOT attempt to use a Class B ping slot until it receives the ***PingSlotChannelAns***.

> **Note:** In order to take advantage of the network-initiated downlink capabilities provided by Class B, the network needs relatively recent information of how to best contact the end-device. This information is provided to the network through any and all uplinks received from the end-device. For this reason it is important for Class B enabled end-devices to send regular uplinks which implicitly inform the network of the best way to contact it as well as providing the network a chance to send new MAC commands which may be required for proper end-device operation.

Octets	3	1
***PingSlotChannelReq* payload**	`Frequency`	`DR`

Table 68: *PingSlotChannelReq* payload format

The `Frequency` coding is identical to the ***NewChannelReq*** MAC command defined for Class A. A value of `0` instructs the end-device that it SHALL use the default frequency plan.

The `DR` octet contains the following fields:

Bits	7:4	3:0
DR	RFU	DataRate

Table 69: DR field format

The `DataRate` subfield is the index of the data rate used for the ping-slot downlinks. The relationship between the index and the physical data rate is defined in [RP002] for each region.

Upon receiving this command, the end-device SHALL answer with a frame containing the ***PingSlotChannelAns*** command. The ***PingSlotChannelAns*** command SHALL be added in the `FOpts` field (if `FPort` is either missing or >`0`) or in the `FRMPayload` field (if `FPort=0`) of all uplinks until a Class A downlink is received by the end-device.

Note: To limit the unavailability of ping slots, the end-device will answer as soon as possible to this MAC command.

Size (octets)	1
***PingSlotChannelAns* payload**	Status

Table 70: *PingSlotChannelAns* payload format

The `Status` bits have the following meaning:

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

Table 71: Status field format

	Bit=0	**Bit=1**
Data rate ok	The designated data rate is not defined for this end-device, the previous data rate is kept	The data rate is compatible with the possibilities of the end-device
Channel frequency ok	The end-device cannot receive on this frequency	This frequency can be used by the end-device

Table 72: Status field bits signification

If either bit equals `0`, the command did not succeed, and the ping-slot parameters SHALL NOT be modified.

12.4 *BeaconTimingReq* and *BeaconTimingAns*

These MAC commands have been deprecated since LoRaWAN L2 1.0.3. End-devices SHALL use ***DeviceTimeReq*** and ***DeviceTimeAns*** commands as substitutes.

13 Class B Beaconing

13.1 Beacon Physical Layer

In addition to relaying frames between end-devices and Network Servers, gateways MAY participate in providing a time-synchronization mechanism by sending beacons at regular fixed intervals. All beacons are transmitted in radio packet implicit mode, that is, without a LoRa physical header and with no CRC appended by the radio.

PHY	`Preamble`	`BCNPayload`

Table 73: Beacon physical format

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

The beacon frame length is tightly coupled to the operation of the radio physical layer. Therefore, the actual frame length and content changes from one region implementation to another. The beacon content, modulation parameters and frequencies are specified in [RP002] for each region.

13.2 Beacon Frame Format

The beacon payload `BCNPayload` consists of a network common part and an OPTIONAL gateway-specific part.

SF 8	Size (octets)		1	4	2	7	3	2
	`BCNPayload`		`Param`	`Time`	CRC	`GwSpecific`	`RFU`	CRC
SF 9	Size (octets)	1	1	4	2	7	2	
	`BCNPayload`	`RFU`	`Param`	`Time`	CRC	`GwSpecific`	CRC	
SF 10	Size (octets)	2	1	4	2	7	1	2
	`BCNPayload`	`RFU`	`Param`	`Time`	CRC	`GwSpecific`	`RFU`	CRC
SF 11	Size (octets)	3	1	4	2	7	2	2
	`BCNPayload`	`RFU`	`Param`	`Time`	CRC	`GwSpecific`	`RFU`	CRC
SF 12	Size (octets)	4	1	4	2	7	3	2
	`BCNPayload`	`RFU`	`Param`	`Time`	CRC	`GwSpecific`	`RFU`	CRC

Table 74: Beacon frame content

The Param bits have the following meaning, where Prec encodes the timing precision of the beacon:

Bits	7:2	1:0
Param	RFU	Prec

Table 75: Param bits

The precision is interpreted as a base-ten exponent of the beacon's transmit time precision $10^{(-6+\mathtt{Prec})}$ **s,** where the default value of Prec is 0. The Prec field can take any value within the range [0:3]. The RFU portion of Param SHALL be set to 0 and the end-device SHALL silently ignore this field. Networks that use a beacon precision value other than 0 for Prec SHOULD send the beacon using a non-default value of BeaconFrequency.

The common part MAY contain an RFU field equal to 0. It also contains a Param including timestamp precision Prec and a timestamp Time expressed in seconds elapsed since January 6, 1980 00:00:00 UTC (start of the GPS epoch) modulo 2^{32}. The integrity of the beacon's network common part is protected by a 16-bit CRC. The value of CRC-16 SHALL be computed on the RFU + Param + Time fields as defined in IEEE 802.15.4-2003, Section 13.4. This CRC SHALL use the polynomial $P(x) = x^{16} + x^{12} + x^5 + x^0$. The CRC SHALL be calculated on the octets in the order they are sent over the air.

13.3 Beacon GwSpecific Field Format

The content of the GwSpecific field is

Size (octets)	1	6
GwSpecific	InfoDesc	Info

Table 76: Beacon GwSpecific field format

The information descriptor InfoDesc describes how the information field Info SHALL be interpreted.

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	NetID + GatewayID
4:127	RFU
128:255	Reserved for custom network-specific broadcasts

Table 77: Beacon InfoDesc index mapping

For a single omnidirectional antenna gateway, the value of `InfoDesc` is `0` when broadcasting GPS coordinates. For a site featuring sector antennas, for example, the first antenna broadcasts the beacon with `InfoDesc=0`, the second antenna with `InfoDesc=1`, and so on.

13.3.1 Gateway GPS coordinate: `InfoDesc=0, 1 or 2`

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

Size (octets)	3	3
Info	Lat	Lng

Table 78: Beacon `Info` field format, `InfoDesc=0,1,2`

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

- The north–south latitude SHALL be encoded using a two's complement 24-bit word, where -2^{23} corresponds to 90° south (the South Pole) and 2^{23} corresponds to 90° north (the North Pole).
- The east–west longitude SHALL be encoded using a two's complement 24-bit word, where -2^{23} corresponds to 180° west and 2^{23} corresponds to 180° east.

> **Note:** It is not possible to describe 90° north because 2^23-1 is the largest number that can be represented in two's complement notation. This approximately 1.2 m position error at the North Pole is considered small.

13.3.2 `NetID` + `GatewayID`

For `InfoDesc=3`, the content of the `Info` field encodes the Network's `NetID` plus a freely allocated gateway or cell identifier. The format of the `Info` field is

Size (octets)	3	3
Info	NetID	GatewayID

Table 79: Beacon `Info` field format, `InfoDesc=3`

13.4 Beacon Encoding Examples

Example: This is a valid EU868 beacon frame (SF9):

`00 00 | 00 00 02 CC | A2 7E | 00 | 01 20 00 | 00 81 03 | DE 55`

Octets are transmitted left to right. The first CRC is calculated on `[00 00 00 00 02 CC]`. The corresponding field values are shown in Table 79: `Beacon Info` field format, `InfoDesc=3`.

Field	RFU	Param	Time	**CRC**	InfoDesc	Lat	Lng	**CRC**
Value Hex	00	00	CC020000	**7EA2**	0	002001	038100	**55DE**

Table 80: Example of beacon CRC calculation (SF9)

The OPTIONAL gateway specific part provides additional information regarding the gateway sending a beacon and therefore can differ for each gateway. The OPTIONAL part is protected by a CRC-16 computed on the `GwSpecific` + `RFU` fields. The CRC-16 definition SHALL be the same as for the mandatory part.

Example: This is a valid US915 beacon (SF12):

Field	RFU	Param	Time	**CRC**	InfoDesc	Lat	Lng	**RFU**	**CRC**
Value Hex	0000	00	CC020000	**7EA2**	00	002001	038100	**00**	**D450**

Table 81: Example of beacon CRC calculation

Over the air, the octets are sent in the following order:

`00 00 00 | 00 00 02 CC | A2 7E | 00 | 01 20 00 | 00 81 03 | 00 | 50 D4`

Listening and synchronizing with the network common part is sufficient to operate a stationary end-device in Class B mode. A mobile end-device MAY also demodulate the gateway-specific part of the beacon to be able to signal to the Network Server whenever it is moving from one cell to another.

13.5 Beaconing Precise Timing

A beacon SHALL be sent every 128s starting on January 6, 1980, 00:00:00 UTC (start of the GPS epoch) plus $T_{\texttt{BeaconDelay}}$. Therefore, a beacon is sent at $B_T = k \times 128 + T_{\texttt{BeaconDelay}} \pm T_{\text{Accuracy}}$ seconds after the GPS epoch, where k is the smallest integer for which $k \times 128 > T$ and T = number of seconds since January 6, 1980 00:00:00 UTC (start of the GPS time).

> **Note:** *T* is GPS time and, unlike Unix time, it increases strictly monotonically and is not influenced by leap seconds.

$T_{\texttt{BeaconDelay}}$ is 1.5 ms. It allows a slight transmit delay (1.5 ms) of the gateways required by their radio system to switch from receive to transmit mode.

The variable T_{Accuracy} denotes the timing precision guaranteed by the gateway. T_{Accuracy} SHALL be provided by the gateway manufacturer with the set of operational conditions required to guarantee it. T_{Accuracy} SHALL be less than or equal to the timing precision indicated by the `Prec` field of the beacon $T_{\text{Accuracy}} \leq 10^{-6+\texttt{Prec}}$ s.

All gateways participating in the broadcast of the Class B beacon must be synchronized. Depending on the timing precision that can be guaranteed by the gateway, two possible modes of the Class B beacon transmission are possible.

$T_{\text{Accuracy}} \leq 1$ µs: gateways tightly synchronized to GPS time.

If the gateway transmissions can be synchronized to the GPS clock with an accuracy of better than 1 µs, the gateway MAY transmit the beacon every 128 s. The `Prec` field of the beacon is set to `0`, indicating a timing accuracy of better than 1 µs for the listening end-devices.

$T_{\text{Accuracy}} > 1$ µs: gateways loosely synchronized to GPS time.

If gateway transmissions can be synchronized to the GPS clock with an accuracy of better than 1 ms, but an accuracy of 1 µs cannot be guaranteed, the transmission of a Class B beacon SHALL be randomized. For each beacon, the gateway SHALL draw a random number *P* with a uniform distribution between 0 and 1. The gateway transmits the beacon if $P < P_{\text{Beacon}}$. If $P \geq P_{\text{Beacon}}$, the gateway remains silent and does not transmit the beacon.

The parameter P_{Beacon} exists on the gateway and MAY be remotely set by the Network Server. The value of P_{Beacon} SHALL be ≤ 0.5 (50%). Different gateways may use different values of the P_{Beacon} parameter.

The parameter *P* may be a pseudo-random number but, in this case, each gateway in the network SHALL use a different seed, resulting in a unique series of *P* values.

If the Class B beacons of two or more loosely synchronized gateways reach the antenna of an end-device with equivalent power, the end-device may not be able to demodulate the beacon. The timing difference between the colliding beacons may be too great to be compensated correctly by the end-device's demodulator. To avoid systematic beacon collision at the antenna of a stationary end-device, randomization must take place. Each loosely synchronized gateway randomly transmits the beacon no more than half of the time. Therefore, although beacon collisions do happen at the end-device's antenna, they are not systematic, and the end-device can still demodulate the Class B beacon with sufficient probability to operate in Class B mode. The P_{Beacon} parameter should be optimized by the network infrastructure based on the average number of gateways that local end-devices can receive. The denser the gateway population becomes, the higher the probability of beacon collision, so the lower the parameter should be. The optimal parameter value depends on too many network-driven factors and is beyond the scope of this specification.

All end-device ping slots SHALL use the start of the receipt of the preamble of the Class B beacon as a timing reference. Therefore, the Network Server SHALL take $T_{\texttt{BeaconDelay}}$ into account when scheduling the Class B downlinks.

13.6 Network Downlink Route Update Requirements

When the Network attempts to communicate with an end-device using a Class B downlink slot, it SHOULD transmit the downlink from the gateway closest to the end-device when the most recent uplink was received. Therefore, the Network Server SHOULD keep track of the approximate position of every Class B end-device.

Whenever a Class B end-device moves and changes cells, it SHALL communicate with the Network Server in order to update its downlink route. This update is performed simply by sending a confirmed or unconfirmed uplink, possibly without applicative payload.

The end-device can communicate in accordance with two basic strategies:

- **Systematic periodic uplink**: This is the simplest method as it does not require demodulation of the gateway-specific field of the beacon. It is applicable only to slowly moving or stationery end-devices. No requirements are imposed on such periodic uplinks.
- **Uplink on cell change**: The end-device can demodulate the optional gateway-specific field of the beacon. It detects that the ID of the gateway broadcasting the beacon it demodulates has changed and sends an uplink. In that case, the end-device SHALL respect a pseudo-random delay within the range of [0s:120s] between the beacon demodulation and the uplink transmission to ensure that the uplinks of multiple Class B end-devices entering or leaving a cell during the same beacon period will not systematically occur at the same time immediately after the beacon broadcast.

Failure to report a cell change can result in a Class B downlink being temporarily not operational.

14 Class B Unicast and Multicast Downlink Channel Frequencies

The Class B downlink channel selection mechanism depends on the way the Class B beacon is broadcast.

14.1 Single-Channel Beacon Transmission

In certain regions (e.g., EU868), a beacon is transmitted on a single channel. In that case, all unicast and multicast Class B downlinks SHALL use a single frequency channel defined by the ***PingSlotChannelReq*** MAC command. The default frequency is defined in [RP002].

14.2 Frequency-Hopping Beacon Transmission

In certain regions (e.g., US902-928 or CN470-510), a Class B beacon SHALL be transmitted following a frequency-hopping pattern.

In certain regions (e.g., CN470-510), the default Class B downlink channel is subject to the definition in the “Regional Parameters” document [RP002].

In other regions with a hopping beacon, by default Class B ping slots SHALL use a channel that is a function of the `Time` field of the previous beacon (see Section 13.2) and the value of `DevAddr`.

$$\text{Class B ping} - \text{slot channel} = \left[\text{DevAddr} + \text{floor}\left(\frac{\text{BeaconTime}}{\text{BeaconPeriod}}\right)\right] \text{ modulo NbChannel}\,,$$

where

- `BeaconTime` is the 32-bit `Time` field of the current beacon period.
- `BeaconPeriod` is the length of the beacon period (defined as 128 s in the specification).
- floor designates rounding to the immediately lower integer value.
- `DevAddr` is the 32-bit network address of the end-device or multicast group.
- `NbChannel` is the number of channels over which the beacon is frequency hopping.

Class B downlinks therefore hop across `NbChannel` channels (identical to the beacon transmit channels) in the ISM band, and all Class B end-devices are equally spread amongst the `NbChannel` downlink channels.

If the ***PingSlotChannelReq*** command with a valid non-zero `Frequency` argument sets the Class B downlink frequency, then all subsequent ping slots SHOULD be opened on this single frequency independently of the previous beacon frequency.

If the ***PingSlotChannelReq*** command with a zero `Frequency` argument is sent, the end-device SHOULD resume the default frequency plan, that is, Class B ping slots hopping across `NbChannel` channels.

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

2251 CLASS C – CONTINUOUSLY LISTENING

15 Continuously Listening End-Device (Class C)

Class C mode is used for applications that have sufficient power available such that they do not need to minimize the time the radio receiver is active as they do in Class A or Class B applications.

Class C-capable end-devices SHALL NOT enable Class B and Class C concurrently. Class A downlink, however, are always available after an end-device uplink.

A Class C-enabled end-device listens as often as possible using a combination of channel/DR parameters referred to as RXC. The end-device SHALL listen on RXC when it is not (a) transmitting or (b) receiving on RX1 or (c) receiving on RX2, according to the Class A definition. To do so, it SHALL open a short window on RXC parameters between the end of the uplink transmission and the beginning of the RX1 reception window. It SHALL open another RXC window between the end of the RX1 window and the beginning of the RX2 window, and it SHALL switch to RXC reception parameters as soon as the RX2 reception window is closed. This final RXC reception window SHALL remain open until the end-device begins to send another packet.

If the end-device is in the process of demodulating a downlink using the RXC parameters when the RX1 or RX2 window should be opened, it SHALL stop the demodulation and switch to the RX1 or RX2 receive window. This applies even when the RXC and RX2 parameters are identical. The purpose of this rule is to achieve a clear separation between downlinks received during RX1 and RX2 windows (called Class A downlinks) and downlinks received during a Class C window (called Class C downlinks). The same frame counter is incremented, whether the downlink uses RX1/RX2 or RXC.

> **Note:** In order to take advantage of the network-initiated downlink capabilities provided by Class C, the network needs relatively recent information of how to best contact the end-device. This information is provided to the network through any and all uplinks received from the end-device. For this reason it is important for Class C enabled end-devices to send regular uplinks which implicitly inform the network of the best way to contact it as well as providing the network a chance to send new MAC commands which may be required for proper end-device operation.

A Class C downlink SHALL NOT transport any MAC command. If an end-device receives a Class C downlink containing a MAC command, either in the `FOpts` field (if `FPort` is either missing or >`0`) or in the `FRMPayload` field (if `FPort=0`), it SHALL silently discard the entire frame.

In case a Class C confirmed downlink is received, the end-device SHALL NOT answer later than a time period equal to CLASS_C_RESP_TIMEOUT * `NbTrans` + RECEIVE_DELAY2 * (`NbTrans`-1) when the end-device sets its uplink ADR bit, and CLASS_C_RESP_TIMEOUT when the end-device unsets its uplink ADR bit. The default value of CLASS_C_RESP_TIMEOUT is 8s. It can be modified in [RP002], it SHALL not be smaller than RETRANSMIT_TIMEOUT plus the maximum time on air of the uplink frame.

> **Note:** The purpose of this timeout is for the Network Server to know how long to wait for a response from the end-device, before it transmits another Class C confirmed downlink. This ensures that the Network Server can process responses without any ambiguity: there may be only a single confirmed Class C downlink pending an acknowledgement.

The uplink frame sent in response MAY be sent up to `NbTrans` times, but no retransmission SHALL occur after the timeout period. It is RECOMMENDED that the end-device transmits each copy of the uplink frame at a random time within CLASS_C_RESP_TIMEOUT.

As this adds a timing requirement compared to responding to Class A downlinks, the end-device may not send an uplink frame within the timeout, for instance if it has a duty cycle limitation. When such an uplink frame is not sent, the end-device SHALL act as if the uplink frame with the ACK bit set was sent.

After sending a Class C confirmed downlink, a network server SHALL NOT send any other confirmed downlink to the end-device until this timeout expires or an uplink frame is received from that end-device.

After sending a Class A confirmed downlink, a network server SHALL NOT send any other confirmed downlink to the end-device until an uplink frame is received from that end-device.

> **Note:** Class C downlink frames of type Unconfirmed Data may be sent without a minimum delay between them.

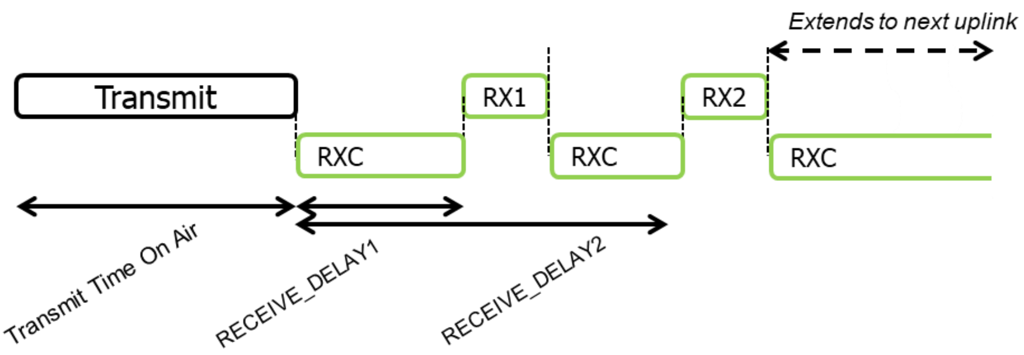

Figure 6: Class C end-device reception slot timing

> **Note:** As defined in Section 3, the end-device does not open RX2 if a downlink is received in RX1. In that case, the end-device opens a continuous RXC receive window at the end of the demodulation of the RX1 downlink that extends until the next uplink.

RXC window parameters

The Class C mechanism may be used to receive unicast or multicast downlink frames:

- Unicast Class C downlinks are typically used for a command and control signal sent to a specific end-device with very low latency (e.g., single streetlight)
- Multicast Class C downlinks are used to broadcast the same downlink frame(s) to a group of end-devices. Typical applications include FUOTA and simultaneously controlling a group of end-evices such as streetlights.

The RXC window parameters differ, depending on whether the Class C functionality is used to receive unicast or multicast downlinks:

- Unicast: The RXC parameters are identical to the RX2 parameters, and they use the same channel and data rate. Modifying the RX2 parameters using the appropriate MAC commands also modifies the RXC parameters.
- Multicast: The RXC parameters are provided by the application layer. All end-devices in the multicast group SHALL share the same RXC parameters. If the multicast RXC parameters are different from the end-device's RX2 parameters, then the end-device is not able to listen simultaneously to multicast and unicast downlink. In that case, the decision whether the end-device should use unicast or multicast RXC parameters is application-specific. If the multicast RXC parameters provided by the application layer match the current RX2 parameters of the end-device, then the end-device receives both unicast and multicast traffic during the RXC windows.

15.1 Class C Multicast Downlinks

Analogously to Class B, Class C end-devices can receive multicast downlink frames. The multicast address and associated network session key SHALL come from the Network Server, and the application session key SHALL come from the Application Server.

More precisely, inside a given end-device, a multicast group is defined by the following parameters called the multicast group context:

1. A 4-octet network address of the multicast group, common to all end-devices of the group.
2. A multicast group-specific session key, different for every multicast group; all end-devices of a given multicast group have the same session keys.
3. A multicast group-specific downlink frame counter.

> **Example:** [TS005] provides an application layer mechanism to set up a multicast group over the air.

The following limitations apply for Class C multicast downlink frames:

- They SHALL NOT carry MAC commands in the `FOpts` field nor in the payload on port 0 because a multicast downlink does not have the same authentication robustness as a unicast frame. The end-device SHALL discard any multicast frame carrying MAC commands.
- The `ACK` bits SHALL be 0 and the F`Type` field SHALL carry the value for `Unconfirmed Data Down`. The end-device SHALL discard the multicast frame otherwise.
- Given that a Class C end-device keeps its receiver active most of the time, the `FPending` bit does not trigger a specific behavior of the end-device and SHALL NOT be used.

SUPPORT INFORMATION

This subsection is informative.

16 Informative Examples

The following examples are illustrations of the LoRaWAN specification for informational purposes only and are not part of the formal specification.

16.1 Uplink Timing Diagram for Unconfirmed Data Frames

The following diagram illustrates the steps executed by an end-device transmitting a single unconfirmed data frame:

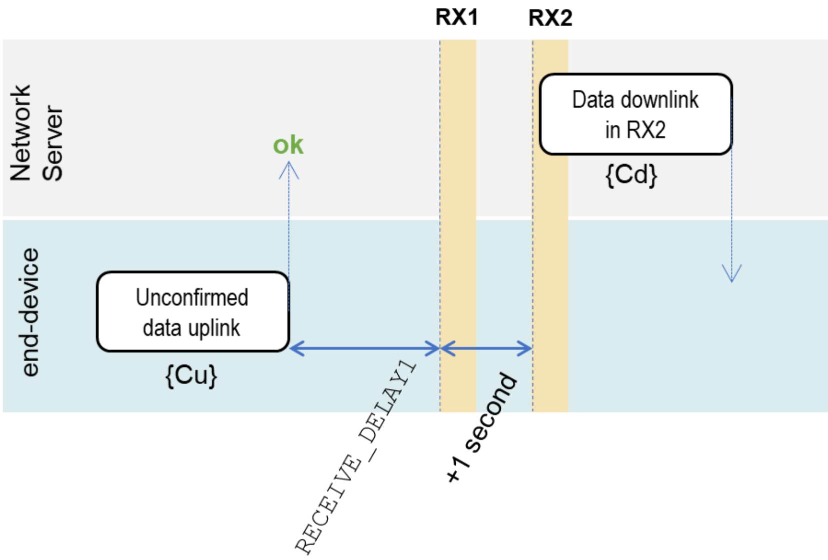

Figure 7: Uplink timing diagram for unconfirmed data frames, `NbTrans=1`

The end-device first transmits an unconfirmed data frame at an arbitrary instant and on an arbitrary channel. The uplink frame counter `Cu` is simply derived by adding 1 to the previous uplink frame counter. The Network Server receives the frame and may generate a downlink frame containing an application payload and/or MAC commands exactly RECEIVE_DELAY1 or RECEIVE_DELAY2 seconds later, using either the first or second Class A receive windows, respectively. This downlink frame uses a data rate and channel specified by the regional channel plan defined in **[RP002]**. The downlink frame counter `Cd` is also derived by adding 1 to the downlink frame counter previously used to transmit to that specific end-device.

The uplink frame counter obeys the constraints imposed by `NbTrans`: the end-device is compelled to transmit the uplink `NbTrans` times or until a downlink is received in the Class A receive windows. In this example, as `NbTrans=1`, the next uplink will be sent with `FCntUp` = `Cu` + 1. This transmission complies with all the specified behaviors defined in this document, including channel selection, timing randomization and duty-cycle limitations.

The following diagram illustrates another example, where an end-device transmits a single payload with unconfirmed uplink data frames, using `NbTrans` = 3.

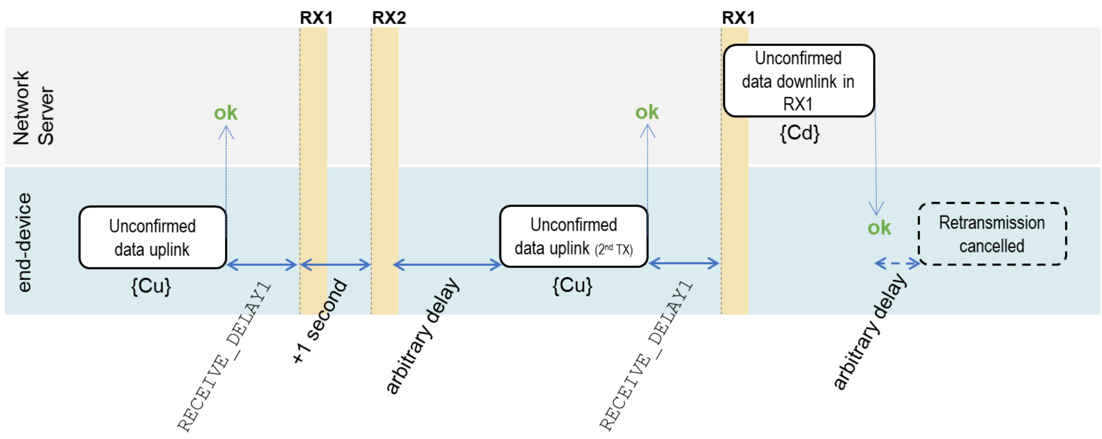

Figure 8: Uplink timing diagram for unconfirmed data frames, `NbTrans`>2

The end-device transmits up to `NbTrans` times the same unconfirmed uplink. The uplink frame counter `Cu` is kept constant for the retransmissions. In this example, first transmission is received correctly by the Network Server. No downlink transmission occurs in the corresponding class A windows, as no response is required: Network Server might not have data or MAC commands to transmit to the end-device, or a downlink transmission to another end-device might be active. In the absence of downlink, end-device waits an arbitrary delay after the end of second class A window, before transmitting the uplink again with same `Cu`, using a different channel. A data downlink frame is then received, on the first class A window following this second transmission. As a consequence of this downlink reception, the third transmission does not occur.

This transmission complies with all the specified behaviors defined in this document, including channel selection, timing randomization and duty-cycle limitations.

16.2 Uplink Timing Diagram for Confirmed Data Frames

The following diagram illustrates the steps executed by an end-device transmitting two confirmed data frames (`Data0` and `Data1`) with `NbTrans=1`.

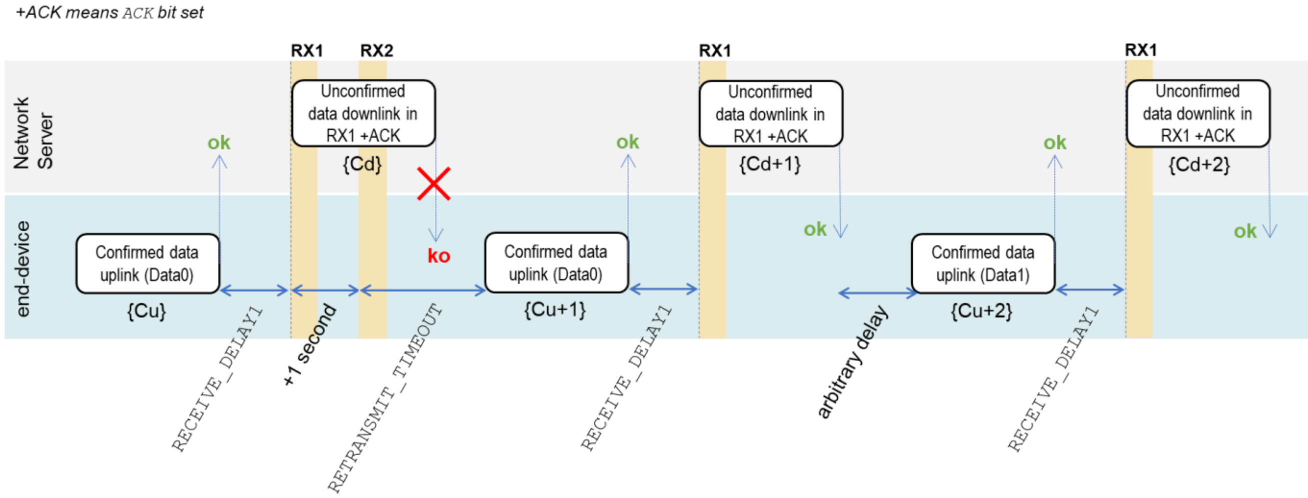

Figure 9: Uplink timing diagram for confirmed data frames

The end-device first transmits a confirmed data frame containing the Data0 payload at an arbitrary instant and on an arbitrary channel. The uplink frame counter Cu is simply derived by adding 1 to the previous uplink frame counter. The Network Server receives the frame and generates a downlink frame with the ACK bit set, which is transmitted exactly RECEIVE_DELAY1 or RECEIVE_DELAY2 seconds later, using the first or second receive window of the end-device, respectively. This downlink frame uses a data rate and channel specified by the regional channel plan defined in [RP002]. The downlink frame counter Cd is also derived by adding 1 to the downlink frame counter previously used to transmit to that specific end-device. If there is no downlink payload pending, the Network Server may send a frame without a payload. In this example, the frame carrying the ACK bit is not received by the end-device. The second receive window is here opened, because frame reception failure happens before this window starts.

If an end-device does not receive a frame with the ACK bit set in one of the two receive windows immediately following the uplink transmission, it may resend the payload after waiting at least RETRANSMIT_TIMEOUT seconds after RX2. The uplink frame counter obeys the constraints imposed by NbTrans: the end-device is compelled to transmit the uplink NbTrans times or until a downlink is received in the Class A receive windows. In this example, as NbTrans=1, the repeated payload will be sent with FCntUp = Cu+1. After this repeated payload, the end-device receives the ACK downlink during its RX1, with FCntDn = Cd+1. The end-device is then free to transmit a new frame on a new channel and is not required to open RX2.

The downlink frames in this example carry an application payload. A downlink frame can carry any combination of ACK, MAC control commands and payload.

16.3 Downlink Diagram for Confirmed Data Frames

The following diagram illustrates the basic sequence of a confirmed downlink data frame.

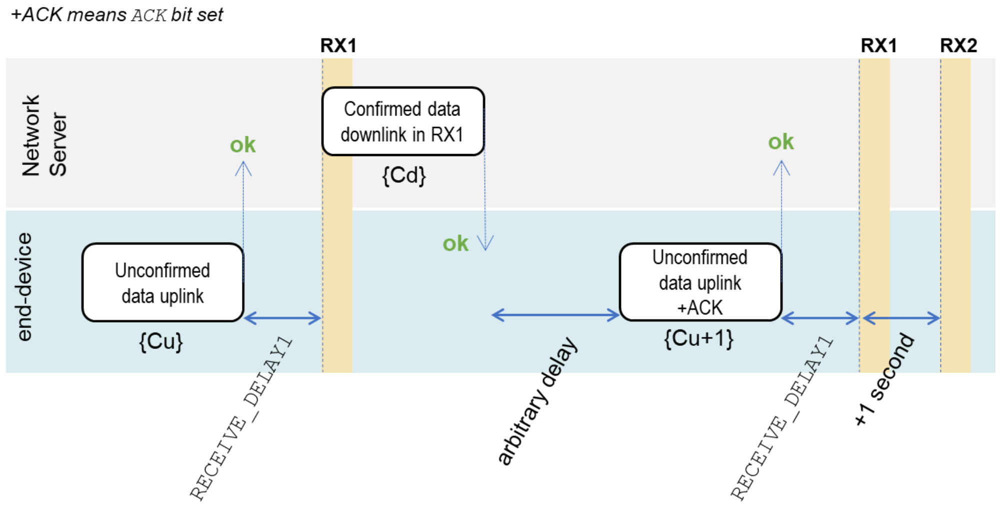

Figure 10: Downlink timing diagram for confirmed data frames

The frame exchange may be initiated by the end-device transmitting a confirmed or unconfirmed data frame, or autonomously by the Network Server using the end-device's Class B ping-slot or Class C RXC window (if the end-device is currently Class B-enabled or Class C-enabled). Upon receiving the downlink data frame requiring an acknowledgement, the end-device transmits an uplink data frame with the `ACK` bit set at the time of its choosing. This frame might also contain piggybacked application payload data and/or MAC commands. This uplink is treated like any other uplink, and as such, this transmission complies with all specified behaviors defined in this document, including channel selection, timing randomization and duty-cycle limitations.

> **Note:** To allow end-devices to be as simple as possible and keep as few states as possible, the end-device may transmit an explicit (possibly empty) acknowledgement data frame immediately after receiving a data frame requiring an acknowledgment. Alternatively, the end-device may defer the transmission of an acknowledgement to piggyback it with its next data frame.
>
> As there is no specified timing of the acknowledgement sent by the end-device, out-of-band coordination between the Network Server and the Class B or Class C-capable end-device is recommended to prevent excessive retransmissions of confirmed downlinks on Class B ping-slots or Class C RXC windows.

16.4 Downlink Timing for Frame-Pending Frames

The following diagram illustrates the use of the `FPending` bit on a downlink. The `FPending` bit can only be set on a downlink frame. When present in a downlink received in a Class A receive window, the `FPending` bit informs the end-device that the Network has one or more downlink frames pending for the end-device. When present in a downlink received in a Class B ping-slot, `FPending` is used to prioritize conflicting ping slots, see Section 9.2 for details. `FPending` has no meaning in a downlink received in a Class C RXC window.

If a Class A end-device receives a frame where `FPending` is set, it is recommended that the end-device transmit an uplink data frame as soon as reasonably possible, which would allow the end-device to receive more information from the Network. However, the precise timing of the uplink transmission is not specified and is application-dependent.

> **Note:** The `FPending` bit is independent of the frame-acknowledgment mechanism.

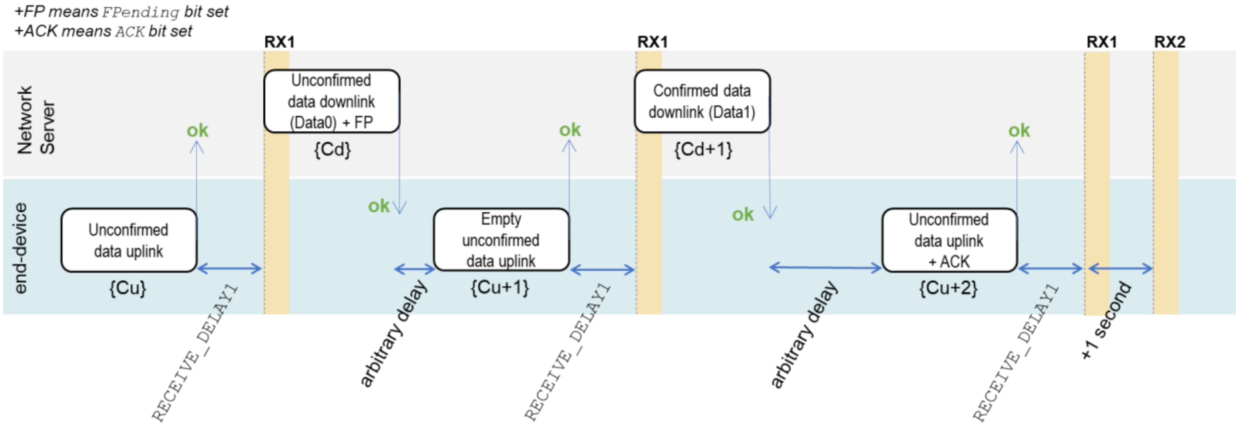

Figure 11: Downlink timing diagram for frame-pending frames, example 1

In this example, the Network Server has two data frames to transmit to the end-device. The frame exchange is initiated by the end-device via a normal unconfirmed uplink data frame. The Network Server uses the first receive window to transmit the `Data0` downlink data frame with the bit `FPending` set, as an unconfirmed downlink data frame. The end-device, in response to the `FPending` indication, transmits quickly an empty unconfirmed data frame. Exactly RECEIVE_DELAY1 seconds later, the Network Server transmits the second downlink data frame `Data1`, using a confirmed downlink data frame but with the `FPending` bit cleared. The end-device transmits an uplink data frame with the `ACK` bit set to acknowledge the confirmed downlink data frame `Data1`.

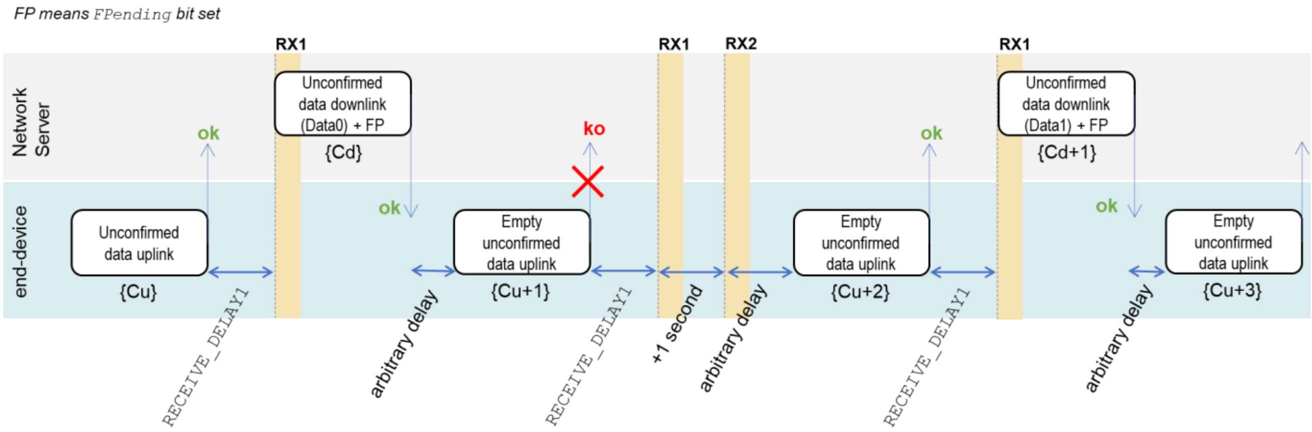

Figure 12: Downlink timing diagram for frame-pending frames, example 2

In this example, the downlink data frames are both unconfirmed frames with the `FPending` bit set, and as such, the end-device does not need to transmit an acknowledgement. After receiving the `Data0` unconfirmed downlink data frame with the `FPending` bit set, the end-device sends an empty data frame at a time of its own choosing. As this uplink is not received by the Network, the Network Server is then still waiting for a spontaneous uplink from the end-device to execute the transfer. The end-device may, at its discretion, offer the Network Server more transmission opportunities by sending a new empty data frame.

The `FPending` bit, the `ACK` bit, and payload data may all be present in the same downlink. The following frame exchange is a perfectly valid example with `NbTrans=1`.

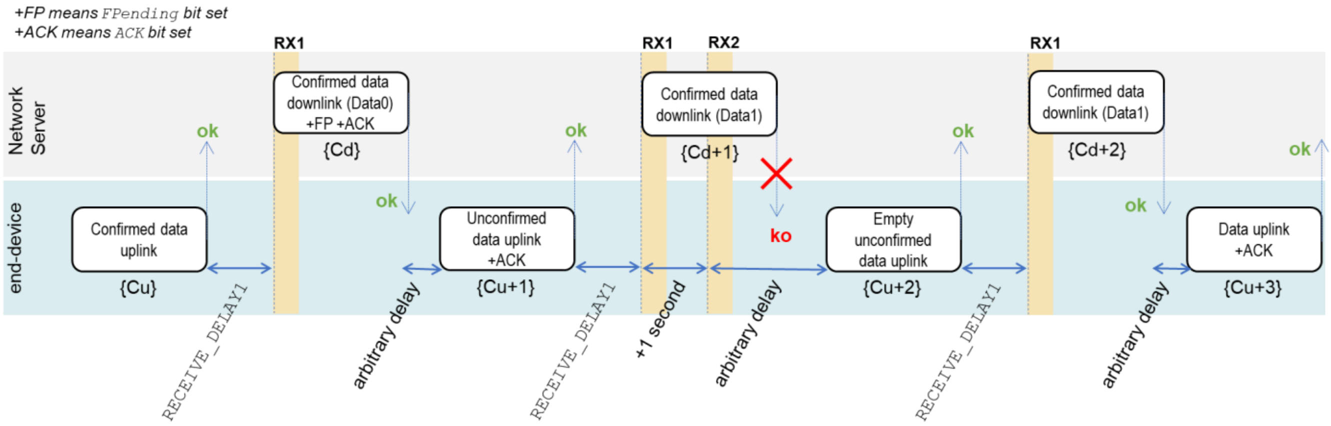

Figure 13: Downlink timing diagram for frame-pending frames, example 3

The end-device sends a confirmed uplink data frame. The Network Server may then answer with a confirmed downlink data frame containing the `Data0` application payload, `ACK` bit set and `FPending` bit set. After an arbitrary delay, the end-device then replies with an uplink with the `ACK` bit set. The end-device misses next downlink containing `Data1`, this data was expected by the end-device because of the previous frame `FPending` indication. After an arbitrary delay, it offers the Network Server another opportunity to send the pending data using an empty frame. This time, `Data1` is received by the end-device, and acknowledged.

17 Revisions

17.1 Revision 1.0

- Approved version of LoRaWAN1.0

17.2 Revision 1.0.1

- Clarified the RX window start time definition
- Corrected the maximum payload size for DR2 in the NA section
- Corrected the typo on the downlink data rate range in 7.2.2
- Introduced a requirement for using coding rate 4/5 in 7.2.2 to guarantee a maximum time on air < 400 mSec
- Corrected the `JoinAccept` MIC calculation in 6.2.5
- Clarified the `NbRep` field and renamed it to `NbTrans` in 5.2
- Removed the possibility to not encrypt the Applicative payload in the MAC layer , removed the paragraph 4.3.3.2. If further security is required by the application , the payload will be encrypted, using any method, at the application layer then re-encrypted at the MAC layer using the specified default LoRaWAN encryption
- Corrected `FHDR` field size typo
- Corrected the channels impacted by `ChMask` when `ChMaskCntl` equals 6 or 7 in 7.2.5
- Clarified 6.2.5 sentence describing the RX1 slot `DataRate` offset in the `JoinResp` frame
- Removed the second half of the `DRoffset` table in 7.2.7 , as `DR>4` will never be used for uplinks by definition
- Removed explicit duty cycle limitation implementation in the EU868 MHz ISM band (Section 7.1)
- Made the ***RXtimingSetupAns*** and ***RXParamSetupAns*** sticky MAC commands to avoid end-device's hidden state problem. (in 5.4 and 5.7)
- Added a frequency plan for the Chinese 470–510 MHz metering band
- Added a frequency plan for the Australian 915–928 MHz ISM band

17.3 Revision 1.0.2

- Extracted Section 7 "Physical layer" that will now be a separate "LoRaWAN regional physical layers definition" document
- corrected the ADR backoff sequence description (ADR_ACK_LIMIT was written instead of ADR_ACK_DELAY) paragraph 4.3.1.1
- Corrected a formatting issue in the title of Section 18.2 (previously Section 19.2 in the 1.0.1 version)
- Added the ***DlChannelRec*** MAC command, this command is used to modify the frequency at which an end-device expects a downlink.
- Added the ***TXParamSetupRec*** MAC command. This command enables to remotely modify the maximum TX dwell time and the maximum radio TX power of an end-device in certain regions
- Added the ability for the end-device to process several ***ADRreq*** commands in a single block in 5.2
- Clarified `AppKey` definition

17.4 Revision 1.0.3

- Imported the Class B chapter from the LoRaWAN1.1 specification
- Added the ***DeviceTimeReq/Ans*** MAC command in the Class A chapter, those commands are required for the Class B beacon acquisition, the MAC commands ***BeaconTimingReq/Ans*** are deprecated.
- Corrected incorrect GPS epoch references
- Corrected various typos

17.5 Revision 1.0.4

- Normative and Grammatical cleanup
- BCP 14 reference added
- Replace `AppEUI` and `AppNonce` with `JoinEUI` and `JoinNonce`
- Clarify Class B and Class C modes of operation as additive to Class A
- Class A RX window opening requirements are clarified
- Reference [RP002] (RP002-1.0.0) as companion document
- Physical-Layer datagrams are referred to as "Packets" as defined in [RP002]
- MAC-Layer datagrams are referred to as "Frames"
- Handling of frames greater than max frame length clarified
- `FPending` clarifications
- Removed MAX_FCNT_GAP
- Clarified `FCnt` usage and behaviors
- FCnts are always 32-bits, and must be persisted by ABP end-devices
- Fports above 224 are not discarded
- Use of "Frame" for MAC-Layer datagrams
- Editorial consistency of Frame type names ("unconfirmed data uplink", etc.)
- Editorial consistency of Join-Request & Join-Accept
- Clarify that "empty frames" are also valid
- ADR behaviors clarified
- Improved the ADR Backoff Example
- Defined Class B bit in `FCtrl` section (instead of just in Class B section)
- MAC Command handling and "sticky" MAC commands overview including priority of responses
- Max, Min and No-Change ***LinkADRReq*** `TXPower`
- Additional ADR Clarifications
- ***LinkCheckAns*** clarifications and definition of `RadioStatus` field as SNR
- Clarify the mandatory nature of MAC commands
- Clarify ***NewChannelReq***/***DlChannelReq*** non-requirement for fixed-channel plan regions
- Require the `DevNonce` to always increments
- Align `DevAddr` `AddrPrefix` definition to Back-End practices
- Require all end-devices to have an associated `DevEUI`, even ABP end-devices
- Clarify channel selection procedure during Join i.e.refer to [RP002]
- Clarify `CFList` handling with respect to other configurable values
- Class C end-devices must successfully uplink once before the network will send downlinks to it
- Retransmission backoff clarified
- Default Ping slot and channel is referred to [RP002]

- Require the support of at least one multicast group
- Define interpretation of `Fpending` in Class B downlinks for unicast and multicast
- ***PingSlotInfoAns*** is defined
- Beacon Frame formats defined for Spreading Factors SF8 to SF12
- Beacon transmission randomization for loosely synched gateways
- Add a time precision field `Prec` to the Beacon to describe the precision of the source gateway's timing and a description was added for its use
- Defined the `Lat`/`Lng` fields for the GPS coordinate fields of the beacon
- Clarified the downlink route update requirements on cell change
- Clarify priority of Class A downlinks over Class C downlinks
- Clarify unicast and multicast RXC parameters and logical model
- Update all of the Informative Examples
- MAC commands from the Network may only be sent on Class A downlinks. Note that regular uplink traffic is expected in class B & C.
- Added a minimum power control range
- For Class B and Class C confirmed downlinks, ACKs shall not occur after a timeout.
- Enforcing duty cycle limit after each uplink frame with Toff
- Modify PingSlotChannelReq acknowledge mechanism (repeated in all uplinks, until next class A downlink, f.k.a. sticky answer)
- After confirmed Class A downlink, NS shall wait for an uplink before sending a Class B or Class C confirmed DL.
- Recommend to send as soon as possible RxParamSetupAns (class C-enabled), and PingSlotChannelAns.

18 Glossary

ABP	Activation By Personalization
ADR	Adaptive Data Rate
AES	Advanced Encryption Standard
CBC	Cipher Block Chaining
CCM	Counter with CBC Message Authentication Code
CMAC	Cipher-based Message Authentication Code
CR	Coding Rate
CRC	Cyclic Redundancy Check
DR	Data Rate
ECB	Electronic Code Book
EIRP	Equivalent Isotopically Radiated Power
ETSI	European Telecommunications Standards Institute
FSK	Frequency Shift Keying modulation technique
HAL	Hardware Abstraction Layer
IP	Internet Protocol
LoRa®	Long Range modulation technique
LoRaWAN®	Long Range Network protocol
MAC	Medium Access Control
MIC	Message Integrity Code
OTAA	Over-The-Air Activation
RF	Radio Frequency
RFU	Reserved for Future Usage
RX	Reception
RSSI	Received Signal Strength Indicator
SF	Spreading Factor
SNR	Signal-to-Noise Ratio
SSL	Secure Socket Layer
TX	Transmission

19 Bibliography

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997

[RFC8174] Leiba, B., “Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words”, BCP 14, RFC8174, May 1997

[IEEE802154]: IEEE Standard for Local and Metropolitan Area Networks—Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs), IEEE Std 802.15.4TM-2011 (Revision of IEEE Std 802.15.4-2006), September 2011.

[RFC4493]: Song, JH., Poovendran, R., Lee, J., and T. Iwata, "The AES-CMAC Algorithm", RFC 4493, RFC Editor, June 2006.

[RP002]: RP002-1.0.0 LoRaWAN Regional Parameters, LoRa Alliance Technical Committee, November 2019

[TS001-1.0.3]: LoRaWAN Link Layer Specification version 1.0.3, LoRa Alliance Technical Committee, July 2018

[TS002]: LoRaWAN Backend Interfaces Specification version 1.0, LoRa Alliance Technical Committee, October 2017

[TS005]: LoRaWAN Remote Multicast Setup Specification v1.0.0, LoRa Alliance Technical Committee, September 2018

[TS009]: LoRaWAN Certification Protocol Specification, LoRa Alliance Certification Committee, 2020.

[TR001]: Technical Recommendations for Preventing State Synchronization Issues around LoRaWAN™ 1.0.x Join Procedure, LoRa Alliance Technical Committee, August 2018.

[NIST-AES] NIST, FIPS 197, "Advanced Encryption Standard (AES)", November 2001.

TS001-1.0.4 LoRaWAN® L2 1.0.4 Specification

NOTICE OF USE AND DISCLOSURE

LoRa Alliance®
5177 Brandin Court
Fremont, CA 94538
United States