

1 **LoRaWAN™ 1.0.3 Specification**

2 Copyright © 2018 LoRa Alliance, Inc. All rights reserved.

3

4 **NOTICE OF USE AND DISCLOSURE**

5 Copyright © LoRa Alliance, Inc. (2017). All Rights Reserved.

6

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

9

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

14

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

21

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

27

28

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

30

31 LoRa Alliance, Inc.
32 3855 SW 153rd Drive
33 Beaverton, OR 97007

34

35 *Note: All Company, brand and product names may be trademarks that are the sole property of their respective owners.*
36
37



39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64

LoRaWAN™ 1.0.3 Specification

Authored by the LoRa Alliance Technical Committee

Chairs:

N.SORNIN (Semtech), A.YEGIN(Actility)

Editor:

N.SORNIN(Semtech)

Contributors:

A.BERTOLAUD (Gemalto), J.CATALANO (Kerlink), J.DELCLEF (ST Microelectronics),
V.DELPORT (Microchip Technology), P.DUFFY (Cisco), F.DYDUCH (Bouygues Telecom),
T.EIRICH (TrackNet), L.FERREIRA (Orange), S.GHAROUT(Orange), O.HERSENT
(Actility), A.KASTTET (Homerider Systems), D.KJENDAL (Senet), V.KLEBAN (Everynet),
J.KNAPP (TrackNet), T.KRAMP (TrackNet), M.KUYPER (TrackNet), P.KWOK (Objenious),
M.LEGOURIEREC (Sagemcom), C.LEVASSEUR (Bouygues Telecom), M.LUIS (Semtech),
M.PAULIAC (Gemalto), P.PIETRI (Orbiwise), D.SMITH (MultiTech), R.SOSS(Actility),
T.TASHIRO (M2B Communications), P.THOMSEN (Orbiwise), A.YEGIN (Actility)

Version: 1.0.3

Date: July 2018

Status: Released

65	Contents	
66	1	Introduction 7
67	1.1	LoRaWAN Classes 7
68	1.2	Conventions 8
69	2	Introduction on LoRaWAN options 9
70	2.1	LoRaWAN Classes 9
71	2.2	Specification scope 10
72	Class A – All end-devices 11	
73	3	Physical Message Formats 12
74	3.1	Uplink Messages 12
75	3.2	Downlink Messages 12
76	3.3	Receive Windows 12
77	3.3.1	First receive window channel, data rate, and start 13
78	3.3.2	Second receive window channel, data rate, and start 13
79	3.3.3	Receive window duration 13
80	3.3.4	Receiver activity during the receive windows 13
81	3.3.5	Network sending a message to an end-device 13
82	3.3.6	Important notice on receive windows 14
83	3.3.7	Receiving or transmitting other protocols 14
84	4	MAC Message Formats 15
85	4.1	MAC Layer (PHYPayload) 15
86	4.2	MAC Header (MHDR field) 15
87	4.2.1	Message type (MType bit field) 16
88	4.2.2	Major version of data message (Major bit field) 16
89	4.3	MAC Payload of Data Messages (MACPayload) 16
90	4.3.1	Frame header (FHDR) 17
91	4.3.2	Port field (FPort) 20
92	4.3.3	MAC Frame Payload Encryption (FRMPayload) 20
93	4.4	Message Integrity Code (MIC) 21
94	5	MAC Commands 22
95	5.1	Link Check commands (<i>LinkCheckReq, LinkCheckAns</i>) 23
96	5.2	Link ADR commands (<i>LinkADRReq, LinkADRAns</i>) 24
97	5.3	End-Device Transmit Duty Cycle (<i>DutyCycleReq, DutyCycleAns</i>) 26
98	5.4	Receive Windows Parameters (<i>RXParamSetupReq, RXParamSetupAns</i>) 26
99	5.5	End-Device Status (<i>DevStatusReq, DevStatusAns</i>) 27
100	5.6	Creation / Modification of a Channel (<i>NewChannelReq, NewChannelAns, DIChannelReq, DIChannelAns</i>) 28
101	5.7	Setting delay between TX and RX (<i>RXTimingSetupReq, RXTimingSetupAns</i>) 30
102	5.8	End-device transmission parameters (<i>TxParamSetupReq, TxParamSetupAns</i>) 30
103	5.9	DeviceTime commands (<i>DeviceTimeReq, DeviceTimeAns</i>) 31
104	6	End-Device Activation 33
105	6.1	Data Stored in the End-device after Activation 33
106	6.1.1	End-device address (DevAddr) 33
107	6.1.2	Application identifier (AppEUI) 33
108	6.1.3	Network session key (NwkSKey) 33
109	6.1.4	Application session key (AppSKey) 33
110	6.2	Over-the-Air Activation 34
111	6.2.1	End-device identifier (DevEUI) 34
112	6.2.2	Application key (AppKey) 34
113	6.2.3	Join procedure 34
114	6.2.4	Join-request message 34
115	6.2.5	Join-accept message 35
116		

117	6.3	Activation by Personalization	37
118	7	Retransmissions back-off	38
119		Class B – Beacon	39
120	8	Introduction to Class B	40
121	9	Principle of synchronous network initiated downlink (Class-B option)	41
122	10	Uplink frame in Class B mode	43
123	11	Downlink Ping frame format (Class B option)	44
124	11.1	Physical frame format.....	44
125	11.2	Unicast & Multicast MAC messages	44
126	11.2.1	Unicast MAC message format.....	44
127	11.2.2	Multicast MAC message format	44
128	12	Beacon acquisition and tracking	45
129	12.1	Minimal beacon-less operation time	45
130	12.2	Extension of beacon-less operation upon reception.....	45
131	12.3	Minimizing timing drift.....	45
132	13	Class B Downlink slot timing	47
133	13.1	Definitions	47
134	13.2	Slot randomization.....	48
135	14	Class B MAC commands	49
136	14.1	PingSlotInfoReq.....	49
137	14.2	BeaconFreqReq	50
138	14.3	PingSlotChannelReq	51
139	14.4	BeaconTimingReq & BeaconTimingAns.....	52
140	15	Beaconing (Class B option).....	53
141	15.1	Beacon physical layer	53
142	15.2	Beacon frame content	53
143	15.3	Beacon <i>GwSpecific</i> field format	54
144	15.3.1	Gateway GPS coordinate:InfoDesc = 0, 1 or 2.....	55
145	15.3.2	NetID+GatewayID.....	55
146	15.4	Beaconing precise timing	55
147	15.5	Network downlink route update requirements	56
148	16	Class B unicast & multicast downlink channel frequencies	57
149	16.1	Single channel beacon transmission	57
150	16.2	Frequency-hopping beacon transmission.....	57
151		Class C – Continuously listening	58
152	17	Class C: Continuously listening end-device.....	59
153	17.1	Second receive window duration for Class C	59
154	17.2	Class C Multicast downlinks.....	60
155		Support information.....	61
156	18	Examples and Application Information.....	62
157	18.1	Uplink Timing Diagram for Confirmed Data Messages	62
158	18.2	Downlink Diagram for Confirmed Data Messages.....	62
159	18.3	Downlink Timing for Frame-Pending Messages	63
160	18.4	Data-Rate Adaptation during Message Retransmissions.....	64
161	19	Recommendation on contract to be provided to the network server by the end-	
162		device provider at the time of provisioning	66
163	20	Recommendation on finding the locally used channels	67
164	21	Revisions	68
165	21.1	Revision 1.0.....	68
166	21.2	Revision 1.0.1	68
167	21.3	Revision 1.0.2.....	68
168	21.4	Revision 1.0.3.....	69
169	22	Glossary.....	70

170	23	Bibliography	71
171	23.1	References	71
172	24	NOTICE OF USE AND DISCLOSURE	72
173			

174 **Tables**

175	Table 1: PHY payload format	15
176	Table 2: MAC message types	16
177	Table 3: Major list.....	16
178	Table 4: FPort list.....	21
179	Table 5: MAC commands	23
180	Table 6: Channel state table.....	24
181	Table 7: LinkADRAns status bits signification	26
182	Table 8: RX2SetupAns status bits signification.....	27
183	Table 9: Battery level decoding	27
184	Table 10: NewChannelAns status bits signification	29
185	Table 11: DICHannelAns status bits signification	30
186	Table 12: Del mapping table.....	30
187	Table 13: Beacon timing	47
188	Table 14 : Class B slot randomization algorithm parameters	48
189	Table 15 : Class B MAC command table	49
190	Table 16 : beacon infoDesc index mapping	54

191

192 **Figures**

193	Figure 1: LoRaWAN Classes.....	9
194	Figure 2: Uplink PHY structure	12
195	Figure 3: Downlink PHY structure	12
196	Figure 4: End-device receive slot timing.	13
197	Figure 5: Radio PHY structure (CRC* is only available on uplink messages)	15
198	Figure 6: PHY payload structure	15
199	Figure 7: MAC payload structure	15
200	Figure 8: Frame header structure.....	15
201	Figure 9: LoRa message format elements	15
202	Figure 11 : DeviceTimeAns payload format	32
203	Figure 12: Beacon reception slot and ping slots	42
204	Figure 13 : class B FCtrl fields.....	43
205	Figure 14 : beacon-less temporary operation	45
206	Figure 15: Beacon timing.....	47
207	Figure 16 : PingSlotInfoReq payload format	49
208	Figure 17 : BeaconFreqReq payload format	50
209	Figure 18 : BeaconFreqAns payload format	50
210	Figure 19 : PingSlotChannelReq payload format.....	51
211	Figure 20 : PingSlotFreqAns payload format	51
212	Figure 21 : beacon physical format	53
213	Figure 22 : beacon frame content.....	53
214	Figure 23 : example of beacon CRC calculation (1)	53
215	Figure 24 : example of beacon CRC calculation (2)	54
216	Figure 25 : beacon GwSpecific field format	54

217	Figure 26 : beacon Info field format , infoDesc = 0,1,2	55
218	Figure 27 : beacon Info field format, infoDesc=3	55
219	Figure 28: Class C end-device reception slot timing.....	60
220	Figure 29: Uplink timing diagram for confirmed data messages	62
221	Figure 30: Downlink timing diagram for confirmed data messages	63
222	Figure 31: Downlink timing diagram for frame-pending messages, example 1	63
223	Figure 32: Downlink timing diagram for frame-pending messages, example 2	64
224	Figure 33: Downlink timing diagram for frame-pending messages, example 3	64
225		

226 1 Introduction

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

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

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

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

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

250 **Note:** Maximum transmit duty-cycle and dwell time per sub-band are
251 region specific and are defined in the LoRaWAN "regional phyLayer
252 definition document

253 1.1 LoRaWAN Classes

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

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

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

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

258 **1.2 Conventions**

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

261 In this document,


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

265 2 Introduction on LoRaWAN options

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

269 2.1 LoRaWAN Classes

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



272
273

Figure 1: LoRaWAN Classes

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

292 2.2 Specification scope

293 This LoRaWAN specification describes the additional functions differentiating an end-device
294 higher Class from one of Class A. A higher Class end-device SHALL also implement all the
295 functionality described in the LoRaWAN Class A specification.

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

300

CLASS A – ALL END-DEVICES

301

All LoRaWAN end-devices **MUST** implement Class A features.

302 3 Physical Message Formats

303 The LoRa terminology distinguishes between uplink and downlink messages.


304 3.1 Uplink Messages

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

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

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

311 *Uplink PHY:*



312 **Figure 2: Uplink PHY structure**

313 3.2 Downlink Messages

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

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

318 *Downlink PHY:*



319 **Figure 3: Downlink PHY structure**

320 3.3 Receive Windows


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

323

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

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

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



324
325

Figure 4: End-device receive slot timing.

326 **3.3.1 First receive window channel, data rate, and start**

327 The first receive window RX1 uses a frequency that is a function of the uplink frequency and
 328 a data rate that is a function of the data rate used for the uplink. RX1 opens
 329 RECEIVE_DELAY1¹ seconds (+/- 20 microseconds) after the end of the uplink modulation.
 330 The relationship between uplink and RX1 slot downlink data rate is region specific and detailed
 331 in the “LoRaWAN regional physical layer specification” document. By default the first receive
 332 window datarate is identical to the datarate of the last uplink.

333 **3.3.2 Second receive window channel, data rate, and start**

334 The second receive window RX2 uses a fixed configurable frequency and data rate and opens
 335 RECEIVE_DELAY2⁷ seconds (+/- 20 microseconds) after the end of the uplink modulation.
 336 The frequency and data rate used can be modified through MAC commands (see Section
 337 5).The default frequency and data rate to use are region specific and detailed in the
 338 “LoRaWAN regional physical layer specification” document

339 **3.3.3 Receive window duration**

340 The length of a receive window MUST be at least the time required by the end-device’s radio
 341 transceiver to effectively detect a downlink preamble.

342 **3.3.4 Receiver activity during the receive windows**

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

348 **3.3.5 Network sending a message to an end-device**

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

¹ RECEIVE_DELAY1 and RECEIVE_DELAY2 are described in Chapter 6.

351 3.3.6 Important notice on receive windows

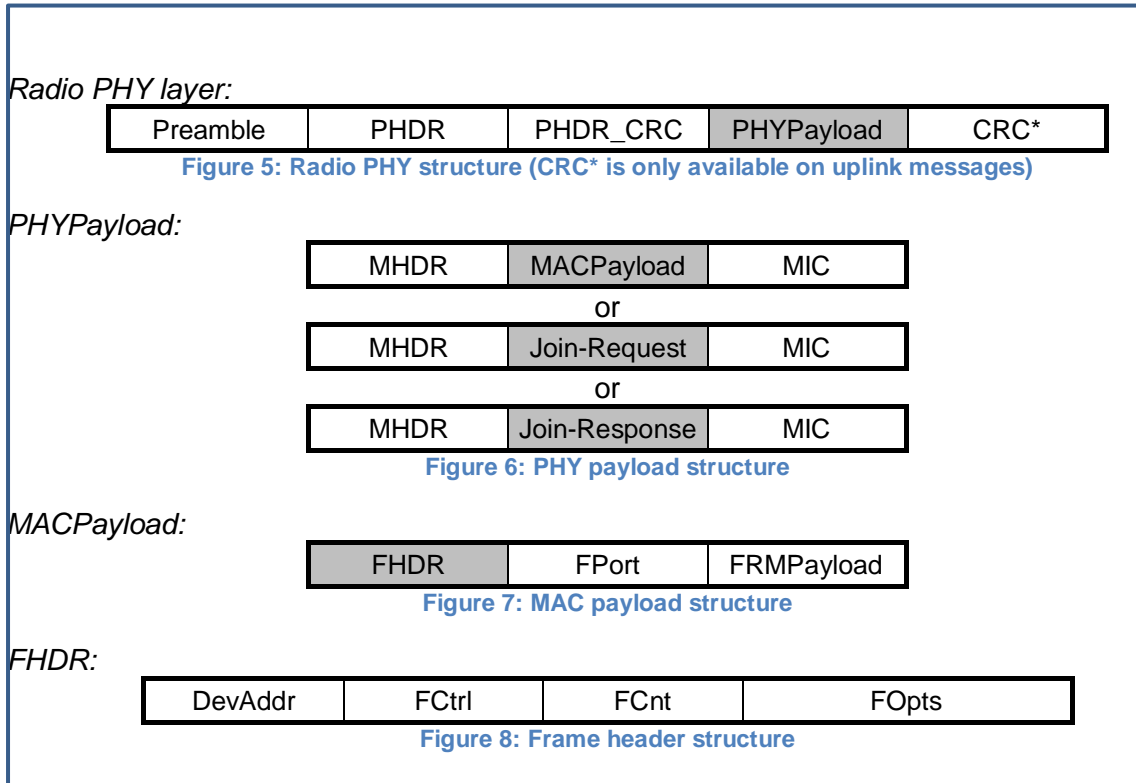
352 An end-device SHALL NOT transmit another uplink message before it either has received a
353 downlink message in the first or second receive window of the previous transmission, or the
354 second receive window of the previous transmission is expired.

355 3.3.7 Receiving or transmitting other protocols

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

359 **4 MAC Message Formats**

360 All LoRa uplink and downlink messages carry a PHY payload (**Payload**) starting with a single-
 361 octet MAC header (**MHDR**), followed by a MAC payload (**MACPayload**)¹, and ending with a
 362 4-octet message integrity code (**MIC**).
 363



374 **Figure 9: LoRa message format elements**

375 **4.1 MAC Layer (PHYPayload)**

376

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

377 **Table 1: PHY payload format**

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

380 **4.2 MAC Header (MHDR field)**

381

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

382

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

¹ Maximum payload size is detailed in the Chapter 6.

385 4.2.1 Message type (MType bit field)

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

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

388 **Table 2: MAC message types**

389 4.2.1.1 Join-request and join-accept messages

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

392 4.2.1.2 Data messages

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

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

401 4.2.2 Major version of data message (Major bit field)

Major bits	Description
00	LoRaWAN R1
01..11	RFU

402 **Table 3: Major list**

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

411 4.3 MAC Payload of Data Messages (MACPayload)

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

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

415 4.3.1 Frame header (FHDR)

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

419

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

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

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

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

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

422 4.3.1.1 Adaptive data rate control in frame header (ADR, ADRACKReq in FCtrl)

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

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

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

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

441 If an end-device whose data rate is optimized by the network to use a data rate higher than
 442 its lowest available data rate, it periodically needs to validate that the network still receives the
 443 uplink frames. Each time the uplink frame counter is incremented (for each new uplink,
 444 repeated transmissions do not increase the counter), the device increments an
 445 ADR_ACK_CNT counter. After ADR_ACK_LIMIT uplinks (ADR_ACK_CNT >=
 446 ADR_ACK_LIMIT) without any downlink response, it sets the ADR acknowledgment request
 447 bit (**ADRACKReq**). The network is required to respond with a downlink frame within the next
 448 ADR_ACK_DELAY frames, any received downlink frame following an uplink frame resets the
 449 ADR_ACK_CNT counter. The downlink **ACK** bit does not need to be set as any response
 450 during the receive slot of the end-device indicates that the gateway has still received the
 451 uplinks from this device. If no reply is received within the next ADR_ACK_DELAY uplinks (i.e.,
 452 after a total of ADR_ACK_LIMIT + ADR_ACK_DELAY), the end-device MAY try to regain
 453 connectivity by switching to the next lower data rate that provides a longer radio range. The
 454 end-device will further lower its data rate step by step every time ADR_ACK_DELAY is

455 reached. The **ADRACKReq** SHALL not be set if the device uses its lowest available data rate
 456 because in that case no action can be taken to improve the link range.

457 **Note:** Not requesting an immediate response to an ADR
 458 acknowledgement request provides flexibility to the network to optimally
 459 schedule its downlinks.

460

461 **Note:** In uplink transmissions the **ADRACKReq** bit is set if
 462 $ADR_ACK_CNT \geq ADR_ACK_LIMIT$ and the current data-rate is
 463 greater than the device defined minimum data rate, it is cleared in
 464 other conditions.

465 4.3.1.2 Message acknowledge bit and acknowledgement procedure (ACK in FCtrl)

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

471 Acknowledgements are only sent in response to the latest message received and are never
 472 retransmitted.

473

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

480 4.3.1.3 Retransmission procedure

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

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

486

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

492

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

499

500

501

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

502 4.3.1.4 Frame pending bit (FPending in FCtrl, downlink only)

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

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

507 4.3.1.5 Frame counter (FCnt)

508 Each end-device has two frame counters to keep track of the number of data frames sent
509 uplink to the network server (FCntUp), incremented by the end-device and received by the
510 end-device downlink from the network server (FCntDown), which is incremented by the
511 network server. The network server tracks the uplink frame counter and generates the
512 downlink counter for each end-device. After a JoinReq – JoinAccept message exchange or a
513 reset for a personalized end-device, the frame counters on the end-device and the frame
514 counters on the network server for that end-device are reset to 0. Subsequently FCntUp and
515 FCntDown are incremented at the sender side by 1 for each new data frame sent in the
516 respective direction. At the receiver side, the corresponding counter is kept in sync with the
517 value received provided the value received has incremented compared to the current counter
518 value and is less than the value specified by MAX_FCNT_GAP¹ after considering counter
519 rollovers. If this difference is greater than the value of MAX_FCNT_GAP then too many data
520 frames have been lost then subsequent will be discarded. The FCnt is not incremented in
521 case of multiple transmissions of an unconfirmed frame (see NbTrans parameter), or in the
522 case of a confirmed frame that is not acknowledged.

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

529 The end-device SHALL not reuse the same FCntUp value, except for retransmission, with the
530 same application and network session keys.

531

532

533

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

534 4.3.1.6 Frame options (FOptsLen in FCtrl, FOpts)

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

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

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

539 If **FOptsLen** is 0, the **FOpts** field is absent. If **FOptsLen** is different from 0, i.e. if MAC
 540 commands are present in the **FOpts** field, the port 0 cannot be used (**FPort** MUST be either
 541 not present or different from 0).

542 MAC commands cannot be simultaneously present in the payload field and the frame options
 543 field. Should this occur, the device SHALL ignore the frame.

544 4.3.2 Port field (FPort)

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

549 Note : The purpose of Fport value 224 is to provide a dedicated Fport to
 550 run Mac compliance test scenarios over-the-air on final versions of
 551 devices, without having to rely on specific test versions of devices for
 552 practical aspects. The test is not supposed to be simultaneous with live
 553 operations, but the Mac layer implementation of the device shall be
 554 exactly the one used for the normal application. The test protocol is
 555 normally encrypted using the AppSKey. This ensures that the network
 556 cannot enable the device's test mode without involving the device's
 557 owner. If the test runs on a live network connected device, the way the
 558 test application on the network side learns the AppSkey is outside of the
 559 scope of the LoRaWAN specification. If the test runs using OTAA on a
 560 dedicated test bench (not a live network), the way the Appkey is
 561 communicated to the test bench, for secured JOIN process, is also
 562 outside of the scope of the specification.

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

565
 566 **FPort** values 225..255 (0xE1..0xFF) are reserved for future standardized application
 567 extensions.
 568

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

569 *N* is the number of octets of the application payload. The valid range for *N* is region specific
 570 and is defined in the “LoRaWAN regional physical layer specification” document.

571 *N* should be equal or smaller than:
 572
$$N \leq M - 1 - (\text{length of FHDR in octets})$$

573 where *M* is the maximum MAC payload length.

574 4.3.3 MAC Frame Payload Encryption (FRMPayload)

575 If a data frame carries a payload, **FRMPayload** MUST be encrypted before the message
 576 integrity code (**MIC**) is calculated.

577 The encryption scheme used is based on the generic algorithm described in IEEE
 578 802.15.4/2006 Annex B [IEEE802154] using AES with a key length of 128 bits.

579 The key *K* used depends on the FPort of the data message:
 580

FPort	K
0	NwkSKey
1..255	AppSKey

581

Table 4: FPort list

582 The fields encrypted are:

583 $pld = \mathbf{FRMPayload}$ 584 For each data message, the algorithm defines a sequence of Blocks A_i for $i = 1..k$ with $k =$
585 $\text{ceil}(\text{len}(pld) / 16)$:

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

586 The direction field (**Dir**) is 0 for uplink frames and 1 for downlink frames.587 The blocks A_i are encrypted to get a sequence S of blocks S_i :

588

589 $S_i = \text{aes128_encrypt}(K, A_i)$ for $i = 1..k$ 590 $S = S_1 | S_2 | .. | S_k$

591 Encryption and decryption of the payload is done by truncating

592

593 $(pld | \text{pad}_{16}) \text{ xor } S$ 594 to the first $\text{len}(pld)$ octets.595 **4.4 Message Integrity Code (MIC)**596 The message integrity code (**MIC**) is calculated over all the fields in the message.

597

598 $msg = \mathbf{MHDR} | \mathbf{FHDR} | \mathbf{FPort} | \mathbf{FRMPayload}$ 599 whereby $\text{len}(msg)$ denotes the length of the message in octets.600 The **MIC** is calculated as follows [RFC4493]:

601

602 $cmac = \text{aes128_cmac}(\mathbf{NwkSKey}, B_0 | msg)$ 603 $\mathbf{MIC} = cmac[0..3]$

604

605 whereby the block B_0 is defined as follows:

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

606

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

608 5 MAC Commands

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

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

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

620 A MAC command consists of a command identifier (**CID**) of 1 octet followed by a possibly
 621 empty command-specific sequence of octets.
 622

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

CID	Command	Transmitted by		Short Description
		End-device	Gateway	
0x0B to 0x0C	<i>RFU</i>			
0x0D	<i>DeviceTimeReq</i>	x		Used by an end-device to request the current date and time
0x0D	<i>DeviceTimeAns</i>		x	Sent by the network, answer to the DeviceTimeReq request
0x0E to 0x7F	<i>RFU</i>			
0x80 to 0xFF	Proprietary	x	x	Reserved for proprietary network command extensions

623

Table 5: MAC commands

624

625

626

627

628

629

630

631

632

633

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

634

635

636

637

638

639

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

640

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

641

642

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

643

644

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

645

646

Size (bytes)	1	1
<i>LinkCheckAns</i> Payload	Margin	GwCnt

647

648

649

650

651

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

652

653

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

654 **5.2 Link ADR commands (*LinkADRReq*, *LinkADRAns*)**
 655

 656 With the *LinkADRReq* command, the network server requests an end-device to perform a
 657 rate adaptation.
 658

Size (bytes)	1	2	1
LinkADRReq Payload	DataRate_TXPower	ChMask	Redundancy

659

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

660

 661 The requested data rate (**DataRate**) and TX output power (**TXPower**) are region-specific and
 662 are encoded as indicated in the “LoRaWAN regional physical layer specification” document.
 663 The TX output power indicated in the command is to be considered the maximum transmit
 664 power the device may operate at. An end-device will acknowledge as successful a command
 665 which specifies a higher transmit power than it is capable of using and should, in that case,
 666 operate at its maximum possible power. The channel mask (**ChMask**) encodes the channels
 667 usable for uplink access as follows with bit 0 corresponding to the LSB:
 668

668

669

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

670

Table 6: Channel state table

 671 A bit in the **ChMask** field set to 1 means that the corresponding channel can be used for uplink
 672 transmissions if this channel allows the data rate currently used by the end-device. A bit set
 673 to 0 means the corresponding channels should be avoided.
 674

674

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

675

 676 In the Redundancy bits the **NbTrans** field is the number of transmissions for each uplink
 677 message. This applies only to “unconfirmed” uplink frames. The default value is 1
 678 corresponding to a single transmission of each frame. The valid range is [1:15]. If **NbTrans**==0
 679 is received the end-device should use the default value. This field can be used by the network
 680 manager to control the redundancy of the node uplinks to obtain a given Quality of Service.
 681 The end-device performs frequency hopping as usual between repeated transmissions, it
 682 does wait after each repetition until the receive windows have expired. . Whenever a downlink
 683 message is received during the RX1 slot window, it SHALL stop any further retransmission of
 684 the same uplink message. For class A devices, a reception in the RX2 slot has the same
 685 effect.
 686

 686 The channel mask control (**ChMaskCntl**) field controls the interpretation of the previously
 687 defined **ChMask** bit mask. It controls the block of 16 channels to which the **ChMask** applies.
 688 It can also be used to globally turn on or off all channels using specific modulation. This field

689 usage is region specific and is defined in the “LoRaWAN regional physical layer specification”
 690 document.

691 The network server MAY include multiple LinkAdrReq commands within a single downlink
 692 message. For the purpose of configuring the end-device channel mask, the end-device will
 693 process all contiguous LinkAdrReq messages, in the order present in the downlink message,
 694 as a single atomic block command. The end-device will accept or reject all Channel Mask
 695 controls in the contiguous block, and provide consistent Channel Mask ACK status indications
 696 for each command in the contiguous block in each LinkAdrAns message, reflecting the
 697 acceptance or rejection of this atomic channel mask setting. The device will only process the
 698 DataRate, TXPower and NbTrans from the last message in the contiguous block, as these
 699 settings govern the end-device global state for these values. The end-device will provide
 700 consistent ACK status in each LinkAdrAns message reflecting the acceptance or rejection of
 701 these final settings.

702

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

705

Size (bytes)	1
LinkADRAns Payload	Status

706

707

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

708

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

710

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

711 **Table 7: LinkADRAns status bits signification**

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

 714 **5.3 End-Device Transmit Duty Cycle (*DutyCycleReq*, *DutyCycleAns*)**

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

 718

Size (bytes)	1
DutyCycleReq Payload	DutyCyclePL

 719

Bits	7:4	3:0
DutyCyclePL	RFU	MaxDCycle

 720
 721
 722
 723 The maximum end-device transmit duty cycle allowed is:

724

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

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

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

 729 **5.4 Receive Windows Parameters (*RXParamSetupReq*,
 730 *RXParamSetupAns*)**

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

 734

Size (bytes)	1	3
RXParamSetupReq Payload	DLsettings	Frequency

 735

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

 736
 737 The **RX1DRoffset** field sets the offset between the uplink data rate and the downlink data
 738 rate used to communicate with the end-device on the first reception slot (RX1). As a default
 739 this offset is 0. The offset is used to take into account maximum power density constraints
 740 for base stations in some regions and to balance the uplink and downlink radio link margins.

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

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

751 The payload contains a single status byte.

Size (bytes)	1
RXParamSetupAns Payload	Status

752
 753 The status (**Status**) bits have the following meaning.

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

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

754
 755 **Table 8: RX2SetupAns status bits signification**

756 If either of the 3 bits is equal to 0, the command did not succeed and the previous parameters
 757 are kept.
 758

759 5.5 End-Device Status (*DevStatusReq*, *DevStatusAns*)

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

Size (bytes)	1	1
DevStatusAns Payload	Battery	Margin

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

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

764 **Table 9: Battery level decoding**

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

Bits	7:6	5:0
Status	RFU	Margin

768 **5.6 Creation / Modification of a Channel (*NewChannelReq*,**
 769 ***NewChannelAns*, *DlChannelReq*, *DlChannelAns*)**
 770

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

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

774

775 The channel index (**ChIndex**) is the index of the channel being created or modified. Depending
 776 on the region and frequency band used, the LoRaWAN specification imposes default channels
 777 which **MUST** be common to all devices and cannot be modified by the ***NewChannelReq***
 778 command (cf. Chapter 6). If the number of default channels is N , the default channels go from
 779 0 to $N-1$, and the acceptable range for **ChIndex** is N to 15. A device **MUST** be able to handle
 780 at least 16 different channel definitions. In certain region the device may have to store more
 781 than 16 channel definitions.
 782

783 The frequency (**Freq**) field is a 24 bits unsigned integer. The actual channel frequency in Hz
 784 is $100 \times \mathbf{Freq}$ whereby values representing frequencies below 100 MHz are reserved for future
 785 use. This allows setting the frequency of a channel anywhere between 100 MHz to 1.67 GHz
 786 in 100 Hz steps. A **Freq** value of 0 disables the channel. The end-device has to check that the
 787 frequency is actually allowed by its radio hardware and return an error otherwise.
 788

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

Bits	7:4	3:0
DrRange	MaxDR	MinDR

791

792 Following the convention defined in Section 5.2 the minimum data rate (**MinDR**) subfield
 793 designate the lowest uplink data rate allowed on this channel. For example 0 designates DR0
 794 / 125 kHz. Similarly, the maximum data rate (**MaxDR**) designates the highest uplink data rate.
 795 For example, $\mathbf{DrRange} = 0x77$ means that only 50 kbps GFSK is allowed on a channel and
 796 $\mathbf{DrRange} = 0x50$ means that DR0 / 125 kHz to DR5 / 125 kHz are supported.

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

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

Size (bytes)	1
NewChannelAns Payload	Status

801

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

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

803

	Bit = 0	Bit = 1
Data rate range ok	The designated data rate range exceeds the ones	The data rate range is compatible with the

Channel frequency ok	currently defined for this end-device	possibilities of the end-device
	The device cannot use this frequency	The device is able to use this frequency.

804

Table 10: NewChannelAns status bits signification

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

807

808 The **DIChannelReq** command allows the network to associate a different downlink frequency
809 to the RX1 slot. This command is applicable for all the physical layer specifications supporting
810 the **NewChannelReq** command (for example EU and China physical layers, but not for US or
811 Australia). In regions where that command is not defined, the device SHALL silently drop it.

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

813

Size (bytes)	1	3
DIChannelReq Payload	ChIndex	Freq

814

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

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

821 If the **DIChannelReq** command is defined in the region where the end-device is operating, the
822 end-device acknowledges the reception of a **DIChannelReq** by sending back a
823 **DIChannelAns** command. The **DIChannelAns** command SHALL be added in the FOpt field
824 of all uplinks until a downlink packet is received by the end-device. This guarantees that even
825 in presence of uplink packet loss, the network is always aware of the downlink frequencies
826 used by the end-device.

827 The payload of this message contains the following information:

Size (bytes)	1
DIChannelAns Payload	Status

828

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

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

830

831

Channel frequency ok	Bit = 0	Bit = 1
	The device cannot use this frequency	The device is able to use this frequency.

Uplink frequency exists

The uplink frequency is not defined for this channel , the downlink frequency can only be set for a channel that already has a valid uplink frequency	The uplink frequency of the channel is valid
---	--

832 **Table 11: DIChannelAns status bits signification**

833 If either of those 2 bits equals 0, the command did not succeed and the RX1 slot channel
834 frequency has not been updated.

835
836

837 **5.7 Setting delay between TX and RX (*RXTimingSetupReq*, 838 *RXTimingSetupAns*)**

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

Size (bytes)	1
RXTimingSetupReq Payload	Settings

842

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

Bits	7:4	3:0
Settings	RFU	Del

844

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

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

846 **Table 12: Del mapping table**

847

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

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

853

854 **5.8 End-device transmission parameters (*TxParamSetupReq*, 855 *TxParamSetupAns*)**

856 This MAC command only needs to be implemented for compliance in certain regulatory
857 regions. Please refer to the “LoRaWAN physical layer definitions” document.

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

861	Size (bytes)	1
862	TxParamSetup payload	EIRP_DwellTime

863 The structure of EIRP_DwellTime field is described below:

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

864

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

868

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

869 The maximum EIRP corresponds to an upper bound on the device's radio transmit power. The
 870 device is not required to transmit at that power , but SHALL never radiate more that this
 871 specified EIRP.

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

Coded Value	Dwell Time
0	No Limit
1	400 ms

874

875 When this MAC command is implemented (region specific), the end-device acknowledges the
 876 TxParamSetupReq command by sending a ***TxParamSetupAns*** command. This
 877 ***TxParamSetupAns*** command doesn't contain any payload.

878 When this MAC command is used in a region where it is not required, the device does not
 879 process it and SHALL NOT transmit an acknowledgement.

880

881

882 5.9 DeviceTime commands (DeviceTimeReq, DeviceTimeAns)

883 This MAC command is only available if the device is activated on a LoRaWAN1.0.3 or later
 884 compatible Network Server. LoRaWAN1.0.2 or earlier servers do not implement this MAC
 885 command.

886 With the ***DeviceTimeReq*** command, an end-device MAY request from the network the current
 887 network time. This allows the device to synchronize its internal clock to the network's clock.

888 This is specifically useful to speed-up the acquisition of the class B beacon. The request has
 889 no payload.

890 With the **DeviceTimeAns** command, the Network Server provides the network time to the end
 891 device. The time provided is the network time captured at the end of the uplink transmission.
 892 The command has a 5 bytes payload defined as follows:

893

Size (bytes)	4	1
DeviceTimeAns Payload	32-bit unsigned integer : Seconds since epoch*	8bits unsigned integer: fractional-second in $\frac{1}{2}^8$ second steps

894

Figure 10 : DeviceTimeAns payload format

895 The time provided by the network MUST have a worst case accuracy of +/-100mSec.

896 The **DeviceTimeAns** command MUST be sent as a class A downlink (i.e. over the RX1/RX2
 897 receive windows of the Class A mode). (*) The GPS epoch (i.e Sunday January the 6th 1980
 898 at 00:00:00 UTC) is used as origin. The “seconds” field is the number of seconds elapsed
 899 since the origin. This field is monotonically increasing by 1 every second. To convert this
 900 field to UTC time, the leap seconds must be taken into account.

901
 902
 903
 904
 905

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

906 **6 End-Device Activation**

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

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

912 **6.1 Data Stored in the End-device after Activation**

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

916 **6.1.1 End-device address (DevAddr)**

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

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

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

923 **6.1.2 Application identifier (AppEUI)**

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

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

927 **6.1.3 Network session key (NwkSKey)**

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

932 **6.1.4 Application session key (AppSKey)**

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

941

942 **6.2 Over-the-Air Activation**

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

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

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

950 **Note:** For over-the-air-activation, end-devices are not personalized with
 951 any kind of network key. Instead, whenever an end-device joins a
 952 network, a network session key specific for that end-device is derived to
 953 encrypt and verify transmissions at the network level. This way, roaming
 954 of end-devices between networks of different providers is facilitated.
 955 Using both a network session key and an application session key further
 956 allows federated network servers in which application data cannot be
 957 read or tampered with by the network provider.

958 **6.2.1 End-device identifier (DevEUI)**

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

961 **6.2.2 Application key (AppKey)**

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

966 **6.2.3 Join procedure**

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

969 **6.2.4 Join-request message**

970 The join procedure is always initiated from the end-device by sending a join-request message.
 971

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

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

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

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

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

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

989
$$cmac = aes128_cmac(AppKey, MHDR | AppEUI | DevEUI | DevNonce)$$

 990
$$MIC = cmac[0..3]$$

992 The join-request message is not encrypted.

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

997 6.2.5 Join-accept message

998 The network server will respond to the **join-request** message with a **join-accept** message if
 999 the end-device is permitted to join a network. The join-accept message is sent like a normal
 1000 downlink but uses delays JOIN_ACCEPT_DELAY1 or JOIN_ACCEPT_DELAY2 (instead of
 1001 RECEIVE_DELAY1 and RECEIVE_DELAY2, respectively). The channel frequency and data
 1002 rate used for these two receive windows are identical to the one used for the RX1 and RX2
 1003 receive windows described in the “receive windows” section of the “LoRaWAN regional
 1004 physical layer specification” document

1005 No response is given to the end-device if the join request is not accepted.

1006 The join-accept message contains an application nonce (**AppNonce**) of 3 octets, a network
 1007 identifier (**NetID**), an end-device address (**DevAddr**), a delay between TX and RX (**RxDelay**)
 1008 and an optional list of channel frequencies (**CFList**) for the network the end-device is joining.
 1009 The CFList option is region specific and is defined in the “LoRaWAN regional physical layer
 1010 specification” document.

1011

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

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

² [RFC4493]

1012 The **AppNonce** is a random value or some form of unique ID provided by the network server
 1013 and used by the end-device to derive the two session keys **NwkSKey** and **AppSKey** as
 1014 follows:¹

1015 $NwkSKey = aes128_encrypt(AppKey, 0x01 | AppNonce | NetID | DevNonce | pad_{16})$

1016 $AppSKey = aes128_encrypt(AppKey, 0x02 | AppNonce | NetID | DevNonce | pad_{16})$

1017 The MIC value for a join-accept message is calculated as follows:²

1018 $cmac = aes128_cmac(AppKey,$

1019 $\quad MHDR | AppNonce | NetID | DevAddr | DLSettings | RxDelay | CFList)$

1020 $MIC = cmac[0..3]$

1021 The join-accept message itself is encrypted with the **AppKey** as follows:

1022 $aes128_decrypt(AppKey, AppNonce | NetID | DevAddr | DLSettings | RxDelay | CFList |$

1023 $MIC)$

1024 **Note:** The network server uses an AES decrypt operation in ECB mode
 1025 to encrypt the join-accept message so that the end-device can use an
 1026 AES encrypt operation to decrypt the message. This way an end-device
 1027 only has to implement AES encrypt but not AES decrypt.

1028

1029 **Note:** Establishing these two session keys allows for a federated
 1030 network server infrastructure in which network operators are not able to
 1031 eavesdrop on application data. In such a setting, the application
 1032 provider must support the network operator in the process of an end-
 1033 device actually joining the network and establishing the NwkSKey for
 1034 the end-device. At the same time the application provider commits to
 1035 the network operator that it will take the charges for any traffic incurred
 1036 by the end-device and retains full control over the AppSKey used for
 1037 protecting its application data.

1038 The format of the **NetID** is as follows: The seven LSB of the **NetID** are called **NwkID** and
 1039 match the seven MSB of the short address of an end-device as described before. Neighboring
 1040 or overlapping networks must have different **NwkIDs**. The remaining 17 MSB can be freely
 1041 chosen by the network operator.

1042 The DLsettings field contains the downlink configuration:

1043

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

1044

1045 The RX1DRoffset field sets the offset between the uplink data rate and the downlink data
 1046 rate used to communicate with the end-device on the first reception slot (RX1). By default
 1047 this offset is 0. The offset is used to take into account maximum power density constraints
 1048 for base stations in some regions and to balance the uplink and downlink radio link margins.

1049 The actual relationship between the uplink and downlink data rate is region specific and
 1050 detailed in the “Physical Layer” section

1051 The delay **RxDelay** follows the same convention as the **Delay** field in the **RXTimingSetupReq**
 1052 command.

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

² [RFC4493]

1053 6.3 Activation by Personalization

1054 Under certain circumstances, end-devices can be activated by personalization. Activation by
1055 personalization directly ties an end-device to a specific network by-passing the **join request**
1056 - **join accept** procedure.

1057 Activating an end-device by personalization means that the **DevAddr** and the two session
1058 keys **NwkSKey** and **AppSKey** are directly stored into the end-device instead of the **DevEUI**,
1059 **AppEUI** and the **AppKey**. The end-device is equipped with the required information for
1060 participating in a specific LoRa network when started.

1061 Each device should have a unique set of NwkSKey and AppSKey. Compromising the keys of
1062 one device shouldn't compromise the security of the communications of other devices. The
1063 process to build those keys should be such that the keys cannot be derived in any way from
1064 publicly available information (like the node address for example).

1065 **7 Retransmissions back-off**

1066
1067 Uplink frames that:

- 1068 • Require an **acknowledgement or an answer** by the network or an application
- 1069 server, and are **retransmitted** by the device if the acknowledgement or answer is not
- 1070 received.

1071 **and**

- 1072 • can be triggered by an **external** event causing **synchronization** across a large
- 1073 (>100) number of devices (power outage, radio jamming, network outage,
- 1074 earthquake...)

1075 can trigger a catastrophic, self-persisting, radio network overload situation.

1076

1077 **Note:** An example of such uplink frame is typically the JoinRequest if the
1078 implementation of a group of end-devices decides to reset the MAC
1079 layer in the case of a network outage.

1080 The whole group of end-device will start broadcasting JoinRequest
1081 uplinks and will only stops when receiving a JoinResponse from the
1082 network.

1083

1084

1085 For those frame retransmissions, the interval between the end of the RX2 slot and the next
1086 uplink retransmission SHALL be random and follow a different sequence for every device (For
1087 example using a pseudo-random generator seeded with the device's address) .The
1088 transmission duty-cycle of such message SHALL respect the local regulation and the following
1089 limits, whichever is more constraining:

1090

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

1091

CLASS B – BEACON

1093 8 Introduction to Class B

1094 This section describes the LoRaWAN Class B layer which is optimized for battery-powered
1095 end-devices that may be either mobile or mounted at a fixed location.

1096 End-devices should implement Class B operation when there is a requirement to open receive
1097 windows at fixed time intervals for the purpose of enabling server initiated downlink messages.

1098 LoRaWAN Class B option adds a synchronized reception window on the end-device.

1099 One of the limitations of LoRaWAN Class A is the Aloha method of sending data from the end-
1100 device; it does not allow for a known reaction time when the customer application or the server
1101 wants to address the end-device. The purpose of Class B is to have an end-device available
1102 for reception at a predictable time, in addition to the reception windows that follows the random
1103 uplink transmission from the end-device of Class A. Class B is achieved by having the gateway
1104 sending a beacon on a regular basis to synchronize all end-devices in the network so that the
1105 end-device can open a short additional reception window (called “ping slot”) at a predictable
1106 time during a periodic time slot.

1107 **Note:** The decision to switch from Class A to Class B comes from the
1108 application layer of the end-device. If this class A to Class B switch
1109 needs to be controlled from the network side, the customer application
1110 must use one of the end-device’s Class A uplinks to send back a
1111 downlink to the application layer, and it needs the application layer on
1112 the end-device to recognize this request – this process is not managed
1113 at the LoRaWAN level.

1114 **9 Principle of synchronous network initiated downlink (Class-B** 1115 **option)**

1116 For a network to support end-devices of Class B, all gateways MAY synchronously broadcast
 1117 a beacon providing a timing reference to the end-devices. Based on this timing reference the
 1118 end-devices can periodically open receive windows, hereafter called “ping slots”, which can
 1119 be used by the network infrastructure to initiate a downlink communication. A network initiated
 1120 downlink using one of these ping slots is called a “ping”. The gateway chosen to initiate this
 1121 downlink communication is selected by the Network Server.. For this reason, if an end-device
 1122 moves and detects a change in the identity advertised in the received beacon, it must send
 1123 an uplink to the Network Server so that the server can update the downlink routing path
 1124 database.

1125 Before a device can operate in Class B mode, the following informations MUST be made
 1126 available to the Network Server out-of-band.

- 1127 • The device’s default ping-slot periodicity
- 1128 • Default Ping-slot data rate
- 1129 • Default Ping-slot channel

1130

1131

1132 All end-devices start and join the network as end-devices of Class A. The end-device
 1133 application can then decide to switch to Class B. This is done through the following process:

- 1134 • The end-device application requests the LoRaWAN layer to switch to Class B mode.
 1135 The LoRaWAN layer in the end-device searches for a beacon and returns either a
 1136 BEACON_LOCKED service primitive to the application if a network beacon was found
 1137 and locked or a BEACON_NOT_FOUND service primitive. To accelerate the beacon
 1138 discovery the LoRaWAN layer MAY use the “DeviceTimeReq” MAC command.

- 1139 • Once in Class B mode, the MAC layer sets to 1 the *Class B* bit of the FCTRL field of
 1140 every uplink frame transmitted. This bit signals to the server that the device has
 1141 switched to Class B. The MAC layer will autonomously schedule a reception slot for
 1142 each beacon and each ping slot. When the beacon reception is successful the end-
 1143 device LoRaWAN layer forwards the beacon content to the application together with
 1144 the measured radio signal strength. The end-device LoRaWAN layer takes into
 1145 account the maximum possible clock drift in the scheduling of the beacon reception
 1146 slot and ping slots. When a downlink is successfully demodulated during a ping slot, it
 1147 is processed similarly to a downlink as described in the LoRaWAN Class A
 1148 specification.

- 1149 • A mobile end-device should periodically inform the Network Server of its location to
 1150 update the downlink route. This is done by transmitting a normal (possibly empty)
 1151 “unconfirmed” or “confirmed” uplink. The end-device LoRaWAN layer will appropriately
 1152 set the *Class B* bit to 1 in the frame’s FCtrl field. Optimally this can be done more
 1153 efficiently if the application detects that the node is moving by analyzing the beacon
 1154 content. In that case the end-device MUST apply a random delay (as defined in Section
 1155 15.5 between the beacon reception and the uplink transmission to avoid systematic
 1156 uplink collisions.

- 1157 • At any time the Network Server MAY change the device’s ping-slot downlink frequency
 1158 or data rate by sending a PingSlotChannelReq MAC command.

- 1159 • The device MAY change the periodicity of its ping-slots at any time. To do so, it MUST
 1160 temporarily stop class B operation (unset class B bit in its uplink frames) and send a

1161 PingSlotInfoReq to the Network Server. Once this command is acknowledged the
 1162 device MAY restart class B operation with the new ping-slot periodicity

- 1163 • If no beacon has been received for a given period (as defined in Section 12.2), the
 1164 synchronization with the network is lost. The MAC layer should inform the application
 1165 layer that it has switched back to Class A. As a consequence the end-device
 1166 LoRaWAN layer stops setting the *Class B* bit in all uplinks and this informs the Network
 1167 Server that the end-device is no longer in Class B mode. The end-device application
 1168 can try to switch back to Class B periodically. This will restart this process starting with
 1169 a beacon search.

1170 The following diagram illustrates the concept of beacon reception slots and ping slots.

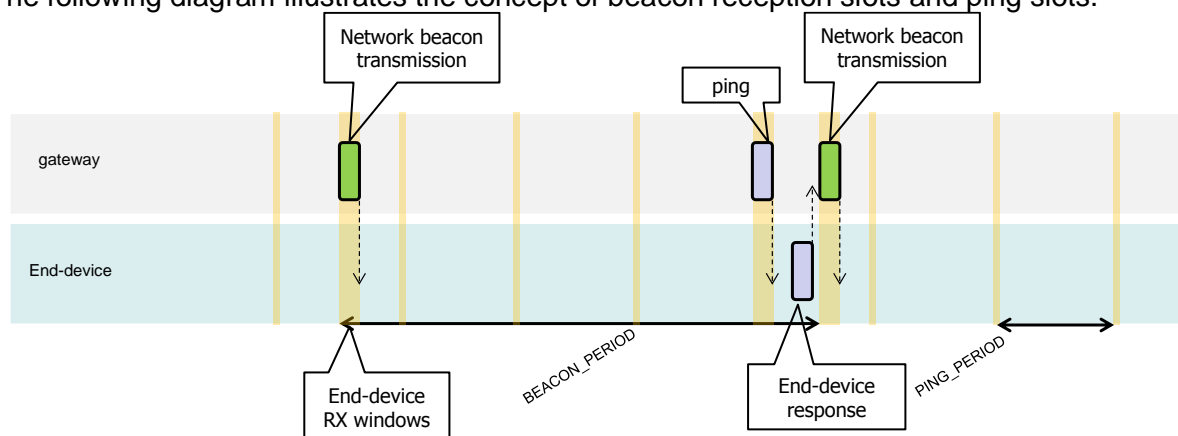


Figure 11: Beacon reception slot and ping slots

1171
 1172

1173 In this example, given the beacon period is 128 s, the end-device also opens a ping reception
 1174 slot every 32 s. Most of the time this ping slot is not used by the server and therefore the end-
 1175 device reception window is closed as soon as the radio transceiver has assessed that no
 1176 preamble is present on the radio channel. If a preamble is detected the radio transceiver will
 1177 stay on until the downlink frame is demodulated. The MAC layer will then process the frame,
 1178 check that its address field matches the end-device address and that the Message Integrity
 1179 Check is valid before forwarding it to the application layer.

1180 **10 Uplink frame in Class B mode**

1181 The uplink frames in Class B mode are same as the Class A uplinks with the exception of
 1182 the RFU bit in the FCtrl field in the Frame header. In the Class A uplink this bit is unused
 1183 (RFU). This bit is used for Class B uplinks.

1184

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

1185

Figure 12 : class B FCtrl fields

1186 The *Class B* bit set to 1 in an uplink signals the Network Server that the device as switched
 1187 to Class B mode and is now ready to receive scheduled downlink pings.

1188

1189 The signification of the FPending bit for downlink is unaltered and still signals that one or
 1190 more downlink frames are queued for this device in the server.

1191

1192 **11 Downlink Ping frame format (Class B option)**

1193 **11.1 Physical frame format**

1194 A downlink Ping uses the same format as a Class A downlink frame but might follow a different
 1195 channel frequency or data rate plan.

1196 **11.2 Unicast & Multicast MAC messages**

1197 Messages can be “unicast” or “multicast”. Unicast messages are sent to a single end-device
 1198 and multicast messages are sent to multiple end-devices. All devices of a multicast group
 1199 **MUST** share the same multicast address and associated encryption keys. The LoRaWAN
 1200 Class B specification does not specify means to remotely setup such a multicast group or
 1201 securely distribute the required multicast key material. This must either be performed during
 1202 the node personalization or through the application layer.

1203 Example: The document [RPD1] describes a possible application layer
 1204 mechanism for over-the-air multicast key distribution.

1205 **11.2.1 Unicast MAC message format**

1206 The MAC payload of a unicast downlink **Ping** uses the format defined in the Class A
 1207 specification. It is processed by the end-device in exactly the same way. The same frame
 1208 counter is used and incremented whether the downlink uses a Class B ping slot or a Class A
 1209 downlink slot.

1210 **11.2.2 Multicast MAC message format**

1211 The Multicast frames share most of the unicast frame format with a few exceptions:

- 1212 • They are not allowed to carry MAC commands, neither in the **FOpt** field, nor in the
 1213 payload on port 0 because a multicast downlink does not have the same authentication
 1214 robustness as a unicast frame.
- 1215 • The **ACK** and **ADRACKReq** bits **MUST** be zero. The **MType** field **MUST** carry the
 1216 value for Unconfirmed Data Down.
- 1217 • The **FPending** bit indicates there is more multicast data to be sent. If it is set the next
 1218 multicast receive slot will carry a data frame. If it is not set the next slot may or may
 1219 not carry data. This bit can be used by end-devices to evaluate priorities for conflicting
 1220 reception slots.

1221

1222 **12 Beacon acquisition and tracking**

1223 Before switching from Class A to Class B, the end-device MUST first receive one of the
1224 network beacons to align his internal timing reference with the network.

1225 Once in Class B, the end-device MUST periodically search and receive a network beacon to
1226 cancel any drift of its internal clock time base, relative to the network timing.

1227 A Class B device may be temporarily unable to receive beacons (out of range from the network
1228 gateways, presence of interference, ..). In this event, the end-device has to gradually widen
1229 its beacon and ping slots reception windows to take into account a possible drift of its internal
1230 clock.

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

1233 **12.1 Minimal beacon-less operation time**

1234 In the event of beacon loss, a device SHALL be capable of maintaining Class B operation for
1235 2 hours (120 minutes) after it received the last beacon. This temporary Class B operation
1236 without beacon is called “beacon-less” operation. It relies on the end-device’s own clock to
1237 keep timing.

1238 During beacon-less operation, Class B unicast, multicast and beacon reception slots MUST
1239 all be progressively expanded to accommodate the end-device’s possible clock drift.
1240

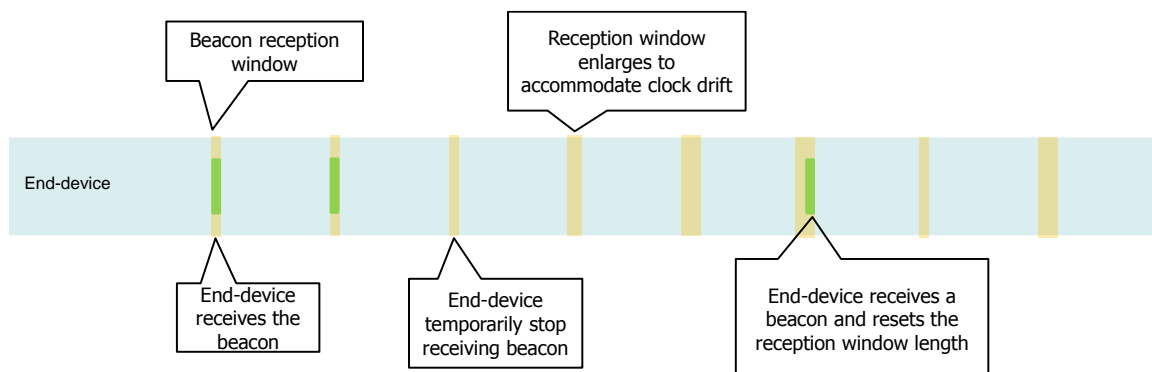


Figure 13 : beacon-less temporary operation

1243 **12.2 Extension of beacon-less operation upon reception**

1244 During this 120 minutes time interval the reception of any beacon directed to the end-device,
1245 should extend the Class B beacon-less operation further by another 120 minutes as it allows
1246 to correct any timing drift and reset the receive slots duration.

Note: Device should also use classB pingSlots downlinks to resynchronize its’ internal clock.

1249 **12.3 Minimizing timing drift**

1250 The end-devices MAY use the beacon’s (when available) precise periodicity to calibrate their
1251 internal clock and therefore reduce the initial clock frequency imprecision. As the timing

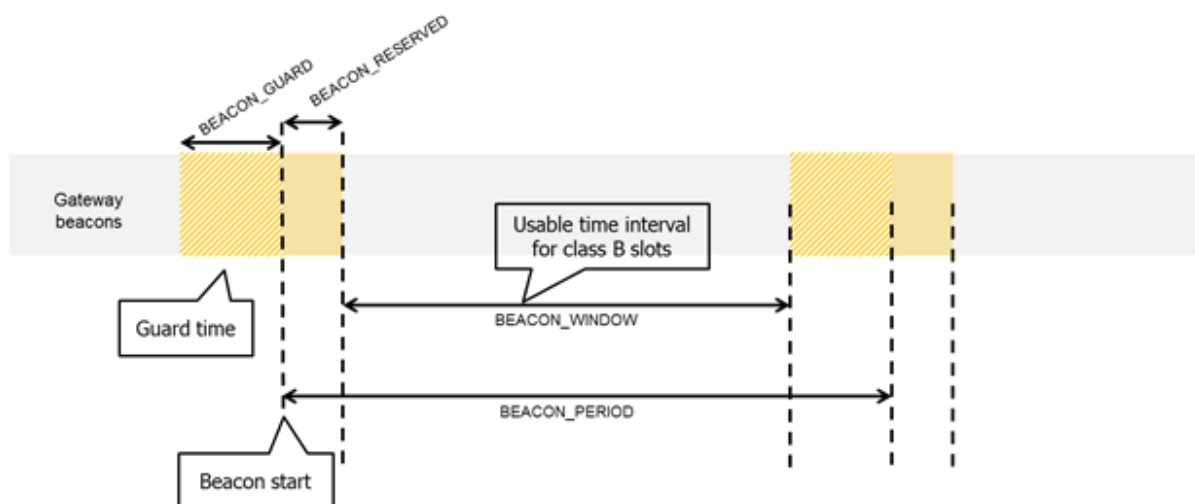
1252 oscillator's exhibit a predictable temperature frequency shift, the use of a temperature sensor
1253 could enable further minimization of the timing drift.

1254 **13 Class B Downlink slot timing**

1255 **13.1 Definitions**

1256 To operate successfully in Class B the end-device MUST open reception slots at precise
1257 instants relative to the infrastructure beacon. This section defines the required timing.

1258 The interval between the start of two successive beacons is called the beacon period. The
1259 beacon frame transmission is aligned with the beginning of the BEACON_RESERVED
1260 interval. Each beacon is preceded by a guard time interval where no ping slot can be placed.
1261 The length of the guard interval corresponds to the time on air of the longest allowed frame.
1262 This is to insure that a Class B downlink initiated during a ping slot just before the guard time
1263 will always have time to complete without colliding with the beacon transmission. The usable
1264 time interval for ping slot therefore spans from the end of the beacon reserved time interval to
1265 the beginning of the next beacon guard interval.



1266
1267

Figure 14: Beacon timing

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

1268

Table 13: Beacon timing

1269 The beacon frame time on air is actually much shorter than the beacon reserved time interval
1270 to allow appending network management broadcast frames in the future.

1271 The beacon window interval is divided into $2^{12} = 4096$ ping slots of 30 ms each numbered from
1272 0 to 4095.

1273 An end-device using the slot number N MUST have its receiver on T_{on} seconds after the start
1274 of the beacon where:

1275
$$T_{on} = beacon_reserved + N * 30 \text{ ms}$$

1276 N is called the *slot index*.

1277 The latest ping slot starts at $beacon_reserved + 4095 * 30 \text{ ms} = 124\,970 \text{ ms}$ after the beacon
1278 start or 3030 ms before the beginning of the next beacon.

1279 **13.2 Slot randomization**

1280 To avoid systematic collisions or over-hearing problems the slot index is randomized and
 1281 changed at every beacon period.

1282 The following parameters are used:

1283

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

1284

Table 14 : Class B slot randomization algorithm parameters

1285 At each beacon period the end-device and the server compute a new pseudo-random offset
 1286 to align the reception slots. An AES encryption with a fixed key of all zeros is used to
 1287 randomize:

1288 $Key = 16 \times 0x00$

1289 $Rand = aes128_encrypt(Key, beaconTime | DevAddr | pad16)$

1290 $pingOffset = (Rand[0] + Rand[1] \times 256) \text{ modulo } pingPeriod$

1291 The slots used for this beacon period will be:

1292 $pingOffset + N \times pingPeriod$ with $N=[0:pingNb-1]$

1293 The node therefore opens receive slots starting at :

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

1294 If the end-device serves simultaneously a unicast and one or more multicast slots this
 1295 computation is performed multiple times at the beginning of a new beacon period. Once for
 1296 the unicast address (the node network address) and once for each multicast group address.

1297 In the case where a multicast ping slot and a unicast ping slot collide and cannot be served
 1298 by the end-device receiver then the end-device should preferentially listen to the multicast
 1299 slot. If there is a collision between multicast reception slots the FPending bit of the previous
 1300 multicast frame can be used to set a preference.

1301 The randomization scheme prevents a systematic collision between unicast and multicast
 1302 slots. If collisions happen during a beacon period then it is unlikely to occur again during the
 1303 next beacon period.

1304 **14 Class B MAC commands**

 1305 All commands described in the Class A specification SHALL be implemented in Class B
 1306 devices. The Class B specification adds the following MAC commands.

1307

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

1308

Table 15 : Class B MAC command table

 1309 **14.1 PingSlotInfoReq**

 1310 With the ***PingSlotInfoReq*** command an end-device informs the server of its unicast ping slot
 1311 periodicity. This command may only be used to inform the server of the periodicity of a
 1312 UNICAST ping slot. A multicast slot is entirely defined by the application and should not use
 1313 this command.

1314

Size (bytes)	1
PingSlotInfoReq Payload	PingSlotParam

1315

Figure 15 : PingSlotInfoReq payload format

Bit#	7:3	[2:0]
PingSlotParam	RFU	Periodicity

 1316 The **Periodicity** subfield is an unsigned 3 bits integer encoding the ping slot period currently
 1317 used by the end-device using the following equation.

1318
$$pingNb = 2^{7-Periodicity} \text{ and } pingPeriod = 2^{5+Periodicity}$$

 1319 The actual ping slot periodicity will be equal to $0.96 \times 2^{Periodicity}$ in seconds

1320 • **Periodicity** = 0 means that the end-device opens a ping slot approximately every
1321 second during the beacon_window interval

1322 • **Periodicity** = 7 , every 128 seconds which is the maximum ping period supported by
1323 the LoRaWAN Class B specification.

1324 To change its ping slot periodicity a device SHALL first revert to Class A , send the new
1325 periodicity through a **PingSlotInfoReq** command and get an acknowledge from the server
1326 through a **PingSlotInfoAns** . It MAY then switch back to Class B with the new periodicity.

1327 This command MAY be concatenated with any other MAC command in the **FHDRFOpt** field
1328 as described in the Class A specification frame format.

1329 14.2 BeaconFreqReq

1330 This command is sent by the server to the end-device to modify the frequency on which this
1331 end-device expects the beacon.

1332

Octets	3
BeaconFreqReq payload	Frequency

1333 **Figure 16 : BeaconFreqReq payload format**

1334 The Frequency coding is identical to the **NewChannelReq** MAC command defined in the
1335 Class A.

1336 **Frequency** is a 24bits unsigned integer. The actual beacon channel frequency in Hz is 100 x
1337 frequ. This allows defining the beacon channel anywhere between 100 MHz to 1.67 GHz by
1338 100 Hz step. The end-device has to check that the frequency is actually allowed by its radio
1339 hardware and return an error otherwise.

1340 A valid non-zero Frequency will force the device to listen to the beacon on a fixed frequency
1341 channel even if the default behavior specifies a frequency hopping beacon (i.e US ISM band).

1342 A value of 0 instructs the end-device to use the default beacon frequency plan as defined in
1343 the “Beacon physical layer” section. Where applicable the device resumes frequency hopping
1344 beacon search.

1345 Upon reception of this command the end-device answers with a **BeaconFreqAns** message.
1346 The MAC payload of this message contains the following information:

Size (bytes)	1
BeaconFreqAns payload	Status

1347 **Figure 17 : BeaconFreqAns payload format**

1348 The **Status** bits have the following meaning:

Bits	7:1	0
Status	RFU	Beacon frequency ok

1349

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

1350

1351 **14.3 PingSlotChannelReq**

 1352 This command is sent by the server to the end-device to modify the frequency and/or the data
 1353 rate on which the end-device expects the downlink pings.

 1354 This command **can only be sent in a class A receive window** (following an uplink). The
 1355 command SHALL NOT be sent in a class B ping-slot. If the device receives it inside a class B
 1356 ping-slot, the MAC command SHALL NOT be processed. Once the NS has sent the first
 1357 PingSlotChannelReq command, it SHOULD resend it until it receives a PingSlotChannelAns
 1358 from the device. It MUST NOT attempt to use a class B ping slot until it receives the
 1359 PingSlotChannelAns.

1360

	Octets	3	1
PingSlotChannelReq Payload	Frequency		DR

1361

Figure 18 : PingSlotChannelReq payload format

 1362 The Frequency coding is identical to the **NewChannelReq** MAC command defined in the
 1363 Class A.

 1364 **Frequency** is a 24bits unsigned integer. The actual ping channel frequency in Hz is 100 x
 1365 frequ. This allows defining the ping channel anywhere between 100MHz to 1.67GHz by 100Hz
 1366 step. The end-device has to check that the frequency is actually allowed by its radio hardware
 1367 and return an error otherwise.

 1368 A value of 0 instructs the end-device to use the default frequency plan. Where applicable the
 1369 device resumes frequency hopping beacon search.

1370 The DR byte contains the following fields:

1371

Bits	7:4	3:0
DR	RFU	data rate

1372

 1373 The “data rate” subfield is the index of the Data Rate used for the ping-slot downlinks. The
 1374 relationship between the index and the physical data rate is defined in [PHY] for each region.

 1375 Upon reception of this command the end-device answers with a **PingSlotFreqAns** message.
 1376 The MAC payload of this message contains the following information:

1377

	Size (bytes)	1
pingSlotFreqAns Payload	Status	

1378

Figure 19 : PingSlotFreqAns payload format

 1379 The **Status** bits have the following meaning:

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

1380

	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 device cannot receive on this frequency	This frequency can be used by the end-device
-----------------------------	---	--

1381
1382
1383
1384
1385

If either of those 2 bits equals 0, the command did not succeed and the ping-slot parameters have not been modified.


1386 **14.4 BeaconTimingReq & BeaconTimingAns**

1387 These MAC commands are deprecated in the LoRaWAN1.0.3 version. The device may use
1388 DeviceTimeReq&Ans commands as a substitute.
1389

1390 15 Beaconing (Class B option)

1391 15.1 Beacon physical layer

1392 Besides relaying messages between end-devices and Network Servers, gateways MAY
 1393 participate in providing a time-synchronization mechanisms by sending beacons at regular
 1394 fixed intervals. All beacons are transmitted in radio packet implicit mode, that is, without a
 1395 LoRa physical header and with no CRC being appended by the radio.
 1396



1397 **Figure 20 : beacon physical format**

1398 The beacon Preamble SHALL begin with (a longer than default) 10 unmodulated symbols.
 1399 This allows end-devices to implement a low power duty-cycled beacon search.

1400 The beacon frame length is tightly coupled to the operation of the radio Physical layer.
 1401 Therefore the actual frame length and content might change from one region implementation
 1402 to another. The beacon content, modulation parameters and frequencies to use are specified
 1403 in [PHY] for each region.

1404 15.2 Beacon frame content

1405 The beacon payload **BCNPayload** consists of a network common part and an OPTIONAL
 1406 gateway-specific part.

Size (bytes)	2 to 5	4	2	7	0 to 3	2
BCNPayload	RFU	Time	CRC	GwSpecific	RFU	CRC
	Compulsory common part			Optional part		

1407 **Figure 21 : beacon frame content**

1408
 1409 The common part contains an RFU field equal to 0, a timestamp **Time** in seconds since
 1410 January 6, 1980 00:00:00 UTC (start of the GPS epoch) modulo 2^{32} . The integrity of the
 1411 beacon's network common part is protected by a 16 bits CRC. The CRC-16 is computed on
 1412 the RFU+Time fields as defined in the IEEE 802.15.4-2003 section 7.2.1.8. This CRC uses
 1413 the following polynomial $P(x) = x^{16} + x^{12} + x^5 + x^0$. The CRC is calculated on the bytes in the
 1414 order they are sent over-the-air

1415 For example: This is a valid EU868 beacon frame:

1416 00 00 | 00 00 02 CC | A2 7E | 00 | 01 20 00 | 00 81 03 | DE 55

1417 Bytes are transmitted left to right. The first CRC is calculated on [00 00 00 02 CC]. The
 1418 corresponding field values are:

Field	RFU	Time	CRC	InfoDesc	lat	long	CRC
Value Hex	0000	CC020000	7EA2	0	002001	038100	55DE

1419 **Figure 22 : example of beacon CRC calculation (1)**

1420
 1421
 1422 The optional gateway specific part provides additional information regarding the gateway
 1423 sending a beacon and therefore may differ for each gateway. The RFU field when applicable

1424 (region specific) should be equal to 0. The optional part is protected by a CRC-16 computed
 1425 on the GwSpecific+RFU fields. The CRC-16 definition is the same as for the mandatory part.

1426 For example: This is a valid US900 beacon:

Field	RFU	Time	CRC	InfoDesc	lat	long	RFU	CRC
Value Hex	000000	CC020000	7E A2	00	002001	038100	00	D450

Figure 23 : example of beacon CRC calculation (2)

1428 Over the air the bytes are sent in the following order:

1429 00 00 00 | 00 00 02 CC | A2 7E | 00 | 01 20 00 | 00 81 03 | 00 | 50 D4

1430 Listening and synchronizing to the network common part is sufficient to operate a stationary
 1431 end-device in Class B mode. A mobile end-device MAY also demodulate the gateway specific
 1432 part of the beacon to be able to signal to the Network Server whenever he is moving from one
 1433 cell to another.

1434 **Note:** As mentioned before, all gateways participating in the beaconing
 1435 process send their beacon simultaneously so that for network common
 1436 part there are no visible on-air collisions for a listening end-device even
 1437 if the end-device simultaneously receives beacons from several
 1438 gateways. Not all gateways are required to participate in the beaconing
 1439 process. The participation of a gateway to a given beacon may be
 1440 randomized. With respect to the gateway specific part, collision occurs
 1441 but an end-device within the proximity of more than one gateway will still
 1442 be able to decode the strongest beacon with high probability.

1443 15.3 Beacon GwSpecific field format

1444 The content of the **GwSpecific** field is as follow:

Size (bytes)	1	6
GwSpecific	InfoDesc	Info

Figure 24 : beacon GwSpecific field format

1446 The information descriptor **InfoDesc** describes how the information field **Info** shall be
 1447 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 16 : beacon infoDesc index mapping

1449
 1450 For a single omnidirectional antenna gateway the **InfoDesc** value is 0 when broadcasting
 1451 GPS coordinates. For a site featuring 3 sectorized antennas for example, the first antenna
 1452 broadcasts the beacon with **InfoDesc** equals 0, the second antenna with **InfoDesc** field
 1453 equals 1, etc.

1454 15.3.1 Gateway GPS coordinate:InfoDesc = 0, 1 or 2

1455 For **InfoDesc** = 0 ,1 or 2, the content of the **Info** field encodes the GPS coordinates of the
1456 antenna broadcasting the beacon

Size (bytes)	3	3
Info	Lat	Lng

1457 **Figure 25 : beacon Info field format , infoDesc = 0,1,2**

1458 The latitude and longitude fields (**Lat** and **Lng**, respectively) encode the geographical location
1459 of the gateway as follows:

- 1460 • The north-south latitude is encoded using a two's complement 24 bit word where -2^{23}
1461 corresponds to 90° south (the South Pole) and $2^{23}-1$ corresponds to $\sim 90^\circ$ north (the
1462 North Pole). The Equator corresponds to 0.
- 1463 • The east-west longitude is encoded using a two's complement 24 bit word where -2^{23}
1464 corresponds to 180° West and $2^{23}-1$ corresponds to $\sim 180^\circ$ East. The Greenwich
1465 meridian corresponds to 0.

1466 15.3.2 NetID+GatewayID

1467 For **InfoDesc** = 3, the content of the **Info** field encodes the network's NetID plus a freely
1468 allocated gateway or cell identifier. The format of the Info field is:

Size (bytes)	3	3
Info	NetID	GatewayID

1469 **Figure 26 : beacon Info field format, infoDesc=3**

1470 15.4 Beaconing precise timing

1471 The beacon is sent every 128 seconds starting at January 6, 1980 00:00:00 UTC (start of the
1472 GPS epoch) plus TBeaconDelay. Therefore the beacon is sent at

$$1473 \quad B_T = k * 128 + TBeaconDelay$$

1474 seconds after the GPS epoch.

1475 whereby k is the smallest integer for which

$$1476 \quad k * 128 > T$$

1477 whereby

1478 T = seconds since January 6, 1980 00:00:00 UTC (start of the GPS time).

1479 **Note:** T is GPS time and unlike Unix time, T is strictly monotonically
1480 increasing and is not influenced by leap seconds.

1481 Whereby TBeaconDelay is 1.5 mSec +/- 1uSec delay.
1482 TBeaconDelay is meant to allow a slight transmission delay of the gateways required by the

1483 radio system to switch from receive to transmit mode.
1484

1485 All end-devices ping slots use the beacon transmission start time as a timing reference,
1486 therefore the Network Server as to take TBeaconDelay into account when scheduling the
1487 class B downlinks.
1488

1489 15.5 Network downlink route update requirements

1490 When the network attempts to communicate with an end-device using a Class B downlink slot,
1491 it transmits the downlink from the gateway which was closest to the end-device when the last
1492 uplink was received. Therefore the Network Server needs to keep track of the rough position
1493 of every Class B device.

1494 Whenever a Class B device moves and changes cell, it needs to communicate with the
1495 Network Server in order to update its downlink route. This update can be performed simply
1496 by sending a “confirmed” or “unconfirmed” uplink, possibly without applicative payload.

1497 The end-device has the choice between 2 basic strategies:

- 1498 • Systematic periodic uplink: simplest method that doesn't require demodulation of the
1499 “gateway specific” field of the beacon. Only applicable to slowly moving or stationery
1500 end-devices. There are no requirements on those periodic uplinks.
- 1501 • Uplink on cell change: The end-device demodulates the optional “gateway specific”
1502 field of the beacon, detects that the ID of the gateway broadcasting the beacon it
1503 demodulates has changed, and sends an uplink. In that case the device SHALL
1504 respect a pseudo random delay in the [0:120] seconds range between the beacon
1505 demodulation and the uplink transmission. This is required to insure that the uplinks
1506 of multiple Class B devices entering or leaving a cell during the same beacon period
1507 will not systematically occur at the same time immediately after the beacon
1508 broadcast.

1509 Failure to report cell change will result in Class B downlink being temporary not operational.
1510 The Network Server may have to wait for the next end-device uplink to transmit downlink
1511 traffic.

1512
1513

1514 **16 Class B unicast & multicast downlink channel frequencies**

1515 The class B downlink channel selection mechanism depends on the way the class B beacon
1516 is being broadcasted.

1517 **16.1 Single channel beacon transmission**

1518 In certain regions (ex EU868) the beacon is transmitted on a single channel. In that case, all
1519 unicast&multicastClass B downlinks use a single frequency channel defined by the
1520 “**PingSlotChannelReq**” MAC command. The default frequency is defined in [PHY].

1521 **16.2 Frequency-hopping beacon transmission**

1522 In certain regions (ex US902-928 or CN470-510) the class B beacon is transmitted following
1523 a frequency hopping pattern.

1524 In that case, by default Class B downlinks use a channel which is a function of the Time field
1525 of the last beacon (see Beacon Frame content) and the DevAddr.

1526 Class B downlink channel = $\left[\text{DevAddr} + \text{floor} \left(\frac{\text{Beacon_Time}}{\text{Beacon_period}} \right) \right] \text{ modulo NbChannel}$

- 1527 • Whereby Beacon_Time is the 32 bit Time field of the current beacon period
- 1528 • Beacon_period is the length of the beacon period (defined as 128sec in the
1529 specification)
- 1530 • Floor designates rounding to the immediately lower integer value
- 1531 • DevAddr is the 32 bits network address of the device
- 1532 • NbChannel is the number of channel over which the beacon is frequency hopping

1533 Class B downlinks therefore hop across NbChannel channels (identical to the beacon
1534 transmission channels) in the ISM band and all Class B end-devices are equally spread
1535 amongst the NbChannel downlink channels.

1536 If the “**PingSlotChannelReq**” command with a valid non-zero argument is used to set the
1537 Class B downlink frequency then all subsequent ping slots should be opened on this single
1538 frequency independently of the last beacon frequency.

1539 If the “**PingSlotChannelReq**” command with a zero argument is sent, the end-device should
1540 resume the default frequency plan, id Class B ping slots hoping across 8 channels.

1541 The underlying idea is to allow network operators to configure end-devices to use a single
1542 proprietary dedicated frequency band for the Class B downlinks if available, and to keep as
1543 much frequency diversity as possible when the ISM band is used.

1544
1545

CLASS C – CONTINUOUSLY LISTENING

1547 **17 Class C: Continuously listening end-device**

1548 The end-devices implementing the Class C option are used for applications that have sufficient
 1549 power available and thus do not need to minimize reception time.

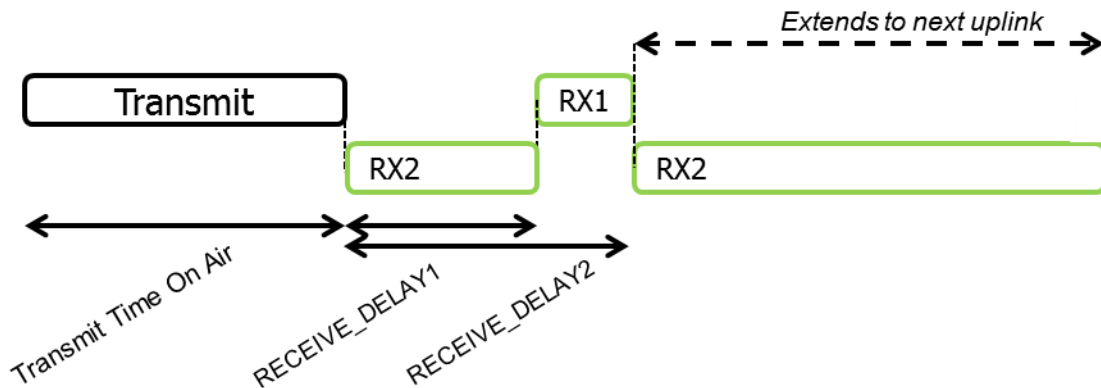
1550 Class C end-devices cannot implement Class B option.

1551 The Class C end-device will listen with RX2 windows parameters as often as possible. The
 1552 end-device listens on RX2 when it is not either (a) sending or (b) receiving on RX1, according
 1553 to Class A definition. To do so, it will open a short window on RX2 parameters between the
 1554 end of the uplink transmission and the beginning of the RX1 reception window and it will
 1555 switch to RX2 reception parameters as soon as the RX1 reception window is closed; the RX2
 1556 reception window will remain open until the end-device has to send another message.


1557 **Note:** There is not specific message for a node to tell the server that it
 1558 is a Class C node. It is up to the application on server side to know that
 1559 it manages Class C nodes based on the contract passed during the join
 1560 procedure.

1561 **17.1 Second receive window duration for Class C**

1562 Class C devices implement the same two receive windows as Class A devices, but they do
 1563 not close RX2 window until they need to send again. Therefore they may receive a downlink
 1564 in the RX2 window at nearly any time, including downlinks sent for the purpose of MAC
 1565 command or ACK transmission. A short listening window on RX2 frequency and data rate is
 1566 also opened between the end of the transmission and the beginning of the RX1 receive
 1567 window.



1568



1569

1570
1571

Figure 27: Class C end-device reception slot timing.

1572 **17.2 Class C Multicast downlinks**

1573 Similarly to Class B, Class C devices may receive multicast downlink frames. The multicast
1574 address and associated network session key and application session key must come from the
1575 application layer.

1576 | Example: [RPD1] provides an application layer mechanism to setup a
1577 | multicast group over-the-air.

1578 The same limitations apply for Class C multicast downlink frames:

- 1579 • They are not allowed to carry MAC commands, neither in the **FOpt** field, nor in the
1580 payload on port 0 because a multicast downlink does not have the same
1581 authentication robustness as a unicast frame.
- 1582 • The **ACK** and **ADRACKReq** bits MUST be zero. The **MType** field MUST carry the
1583 value for Unconfirmed Data Down.
- 1584 • The **FPending** bit indicates there is more multicast data to be sent. Given that a Class
1585 C device keeps its receiver active most of the time, the **FPending** bit does not trigger
1586 any specific behavior of the end-device.

SUPPORT INFORMATION

1587

1588 This sub-section is only a recommendation.

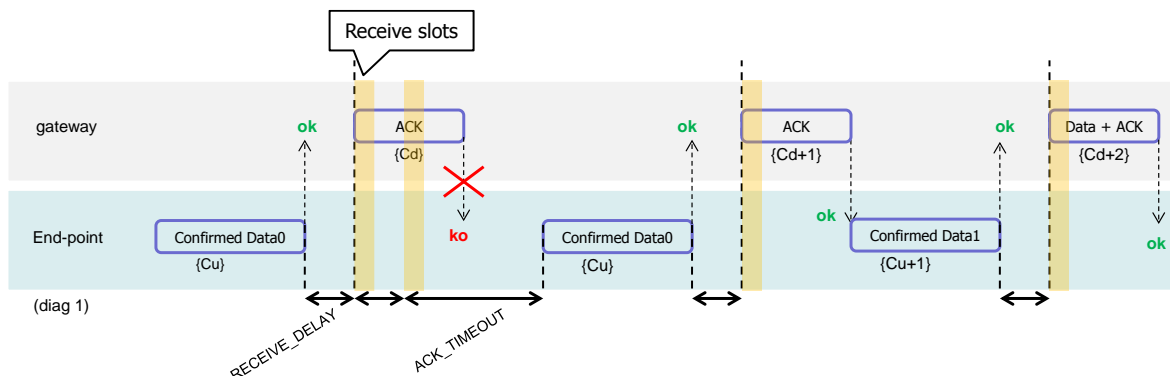
1589

1590 **18 Examples and Application Information**

1591 Examples are illustrations of the LoRaWAN spec for information, but they are not part of the
 1592 formal specification.

1593 **18.1 Uplink Timing Diagram for Confirmed Data Messages**

1594 The following diagram illustrates the steps followed by an end-device trying to transmit two
 1595 confirmed data frames (Data0 and Data1):
 1596



1597 **Figure 28: Uplink timing diagram for confirmed data messages**
 1598

1599 The end-device first transmits a confirmed data frame containing the Data0 payload at an
 1600 arbitrary instant and on an arbitrary channel. The frame counter Cu is simply derived by adding
 1601 1 to the previous uplink frame counter. The network receives the frame and generates a
 1602 downlink frame with the ACK bit set exactly RECEIVE_DELAY1 seconds later, using the first
 1603 receive window of the end-device. This downlink frame uses the same data rate and the same
 1604 channel as the Data0 uplink. The downlink frame counter Cd is also derived by adding 1 to
 1605 the last downlink towards that specific end-device. If there is no downlink payload pending the
 1606 network SHALL generate a frame without a payload. In this example the frame carrying the
 1607 ACK bit is not received.

1608 If an end-device does not receive a frame with the ACK bit set in one of the two receive
 1609 windows immediately following the uplink transmission it MAY resend the same frame with the
 1610 same payload and frame counter again at least ACK_TIMEOUT seconds after the second
 1611 reception window. This resend MUST be done on another channel and MUST obey the duty
 1612 cycle limitation as any other normal transmission. If this time the end-device receives the ACK
 1613 downlink during its first receive window, as soon as the ACK frame is demodulated, the end-
 1614 device is free to transmit a new frame on a new channel.

1615 The third ACK frame in this example also carries an application payload. A downlink frame
 1616 can carry any combination of ACK, MAC control commands and payload.

1617 **18.2 Downlink Diagram for Confirmed Data Messages**

1618 The following diagram illustrates the basic sequence of a “confirmed” downlink.
 1619
 1620

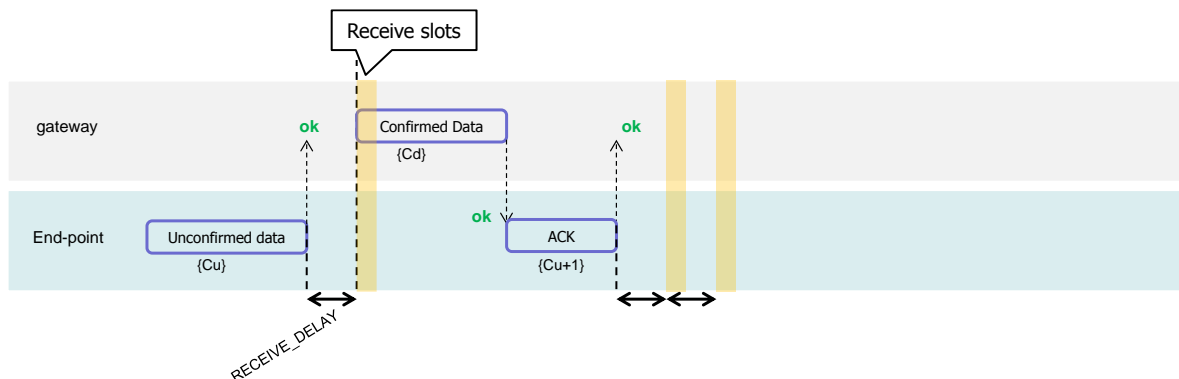


Figure 29: Downlink timing diagram for confirmed data messages

1621
1622

1623 The frame exchange is initiated by the end-device transmitting an “unconfirmed” application
 1624 payload or any other frame on channel A. The network uses the downlink receive window to
 1625 transmit a “confirmed” data frame towards the end-device on the same channel A. Upon
 1626 reception of this data frame requiring an acknowledgement, the end-device transmits a frame
 1627 with the ACK bit set at its own discretion. This frame might also contain piggybacked data or
 1628 MAC commands as its payload. This ACK uplink is treated like any standard uplink, and as
 1629 such is transmitted on a random channel that might be different from channel A.

1630
1631
1632
1633
1634
1635

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

1636 18.3 Downlink Timing for Frame-Pending Messages

1637 The next diagram illustrates the use of the **frame pending** (FPending) bit on a downlink. The
 1638 FPending bit can only be set on a downlink frame and informs the end-device that the network
 1639 has several frames pending for him; the bit is ignored for all uplink frames.

1640 If a frame with the FPending bit set requires an acknowledgement, the end-device shall do so
 1641 as described before. If no acknowledgment is required, the end-device MAY send an empty
 1642 data message to open additional receive windows at its own discretion, or wait until it has
 1643 some data to transmit itself and open receive windows as usual.

1644

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

(*) F_P means 'frame pending' bit set

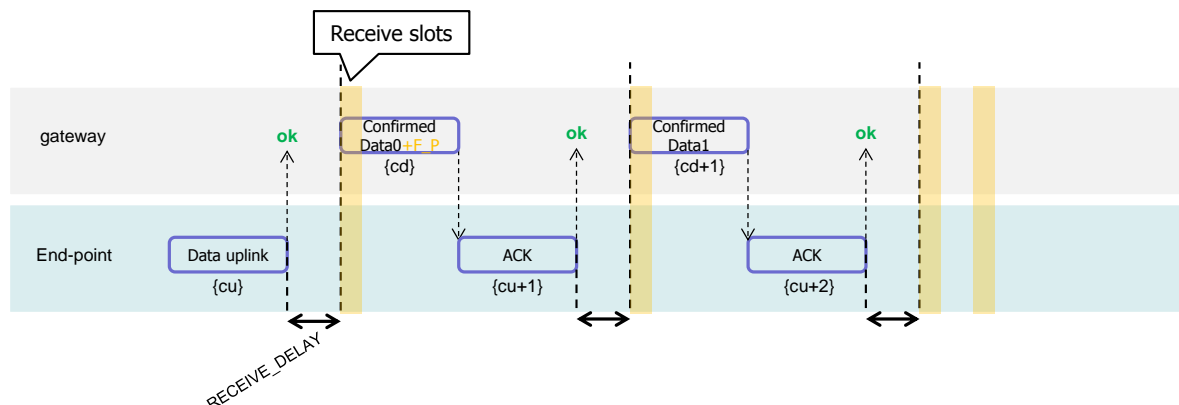
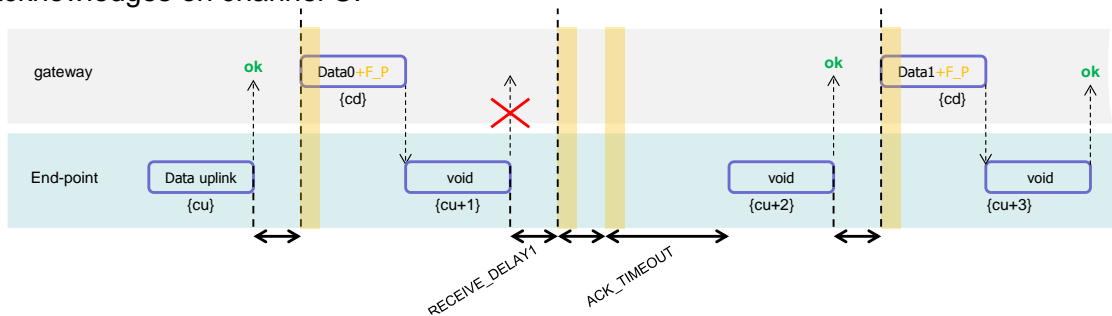


Figure 30: Downlink timing diagram for frame-pending messages, example 1

1645
1646

1647 In this example the network has two confirmed data frames to transmit to the end-device. The
 1648 frame exchange is initiated by the end-device via a normal “unconfirmed” uplink message on
 1649 channel A. The network uses the first receive window to transmit the Data0 with the bit
 1650 FPending set as a confirmed data message. The device acknowledges the reception of the
 1651 frame by transmitting back an empty frame with the ACK bit set on a new channel B.
 1652 RECEIVE_DELAY1 seconds later, the network transmits the second frame Data1 on channel
 1653 B, again using a confirmed data message but with the FPending bit cleared. The end-device
 1654 acknowledges on channel C.

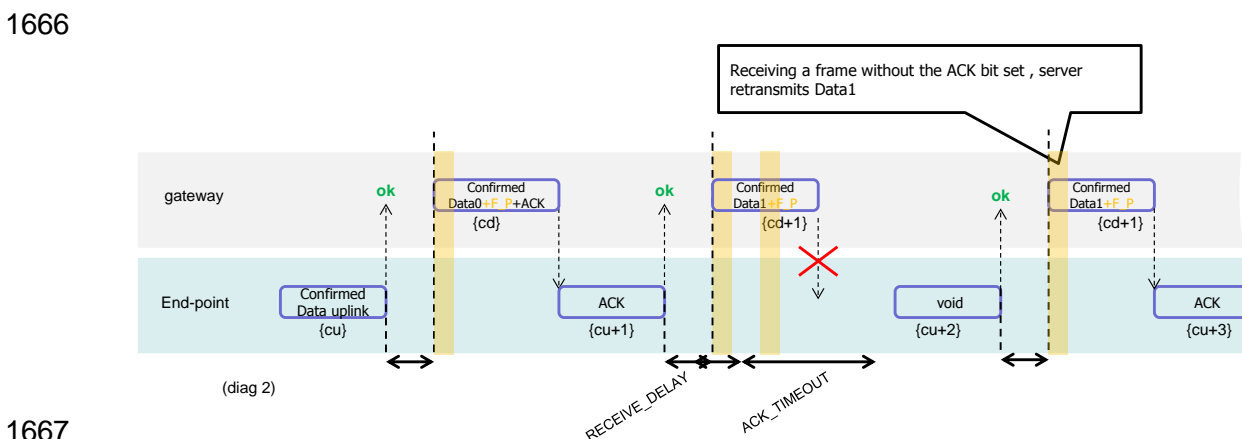


1655
 1656 **Figure 31: Downlink timing diagram for frame-pending messages, example 2**

1657 In this example, the downlink frames are “unconfirmed” frames, the end-device does not need
 1658 to send back and acknowledge. Receiving the Data0 unconfirmed frame with the FPending
 1659 bit set the end-device sends an empty data frame. This first uplink is not received by the
 1660 network. If no downlink is received during the two receive windows, the network has to wait
 1661 for the next spontaneous uplink of the end-device to retry the transfer. The end-device can
 1662 speed up the procedure by sending a new empty data frame.

1663 **Note:** An acknowledgement is never sent twice.

1664 The FPending bit, the ACK bit, and payload data can all be present in the same downlink. For
 1665 example, the following frame exchange is perfectly valid.



1667
 1668 **Figure 32: Downlink timing diagram for frame-pending messages, example 3**

1669 The end-device sends a “confirmed data” uplink. The network can answer with a confirmed
 1670 downlink containing Data + ACK + “Frame pending” then the exchange continues as
 1671 previously described.

18.4 Data-Rate Adaptation during Message Retransmissions

1672 When an end-device attempts the transmission of a “confirmed” frame toward the network it
 1673 expects to receive an acknowledgement in one of the subsequent reception slot. In the
 1674 absence of the acknowledgement it will try to re-transmit the same data again. This re-
 1675

1676 transmission happens on a new frequency channel, but can also happen at a different data
 1677 rate (preferable lower) than the previous one. It is strongly recommended to adopt the
 1678 following re-transmission strategy.

1679 The first transmission of the “confirmed” frame happens with a data rate DR.
 1680

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

1681 The Data Rate $\max(a,b)$ stands for maximum of a and b values.

1682 If after a recommended 8 transmissions, the frame has not been acknowledged the MAC layer
 1683 should return an error code to the application layer.

1684 **Note:** For each re-transmission, the frequency channel is selected
 1685 randomly as for normal transmissions.

1686 Any further transmission uses the last data rate used.

1687 For example if an end-device sends a “confirmed” frame first using DR5 and has to retransmit
 1688 3 times (twice at DR5 then twice at DR4), the next frame transmitted will use DR4

1689 Other example, if an end-device sends a “confirmed” frame first using DR5 and does not
 1690 receive an acknowledge after 8 transmissions (2 at DR5, 2 at DR4, ... , 2 at DR2), and the
 1691 application of this end-device re-initiates a “confirmed” transmission a little later, the first two
 1692 transmission will be tentatively at DR2, then switch to DR1, then to DR0.

1693 **19 Recommendation on contract to be provided to the network**
1694 **server by the end-device provider at the time of provisioning**

1695 Configuration data related to the end-device and its characteristics must be known by the
1696 network server at the time of provisioning. –This provisioned data is called the “contract”. This
1697 contract cannot be provided by the end-device and must be supplied by the end-device
1698 provider using another channel (out-of-band communication).

1699 This end-device contract is stored in the network server. It can be used by the application
1700 server and the network controller to adapt the algorithms.

1701 This data will include:

- 1702 • End-device specific radio parameters (device frequency range, device maximal
1703 output power, device communication settings - RECEIVE_DELAY1,
1704 RECEIVE_DELAY2)
- 1705 • Application type (Alarm, Metering, Asset Tracking, Supervision, Network Control)

1706 20 Recommendation on finding the locally used channels

1707 End-devices that can be activated in territories that are using different frequencies for
1708 LoRaWAN will have to identify what frequencies are supported for join message at their
1709 current location before they send any message. The following methods are proposed:

- 1710 • A GPS enabled end-device can use its GPS location to identify which frequency band
1711 to use.
- 1712 • End-device can search for a beacon and use its frequency to identify its region
- 1713 • End-device can search for a beacon and if this one is sending the antenna GPS
1714 coordinate, it can use this to identify its region
- 1715 • End-device can search for a beacon and if this one is sending a list of join frequencies,
1716 it can use this to send its join message

1717 **21 Revisions**

 1718 **21.1 Revision 1.0**

- 1719
- Approved version of LoRaWAN1.0

 1720 **21.2 Revision 1.0.1**

- 1721
- 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 < 400mSec
 - 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 data rate offset in the JoinResp message
 - 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 EU868Mhz ISM band (chapter7.1)
 - 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 470-510MHz metering band
 - Added a frequency plan for the Australian 915-928MHz ISM band

 1745 **21.3 Revision 1.0.2**

- 1746
- Extracted section 7 “Physical layer” that will now be a separated document “LoRaWAN regional physical layers definition”
 - corrected the ADR_backoff sequence description (ADR_ACK_LIMT 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 DIChannelRec MAC command, this command is used to modify the frequency at which an end-device expects a downlink.
 - Added the Tx ParamSetupRec MAC command. This command enables to remotely modify the maximum TX dwell time and the maximum radio transmit power of a 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
- 1759
-
- 1760

1761 21.4 Revision 1.0.3

1762

1763

1764

1765

1766

1767

1768

1769

1770

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

1771 22 Glossary

1772		
1773	ADR	Adaptive Data Rate
1774	AES	Advanced Encryption Standard
1775	AFA	Adaptive Frequency Agility
1776	AR	Acknowledgement Request
1777	CBC	Cipher Block Chaining
1778	CMAC	Cipher-based Message Authentication Code
1779	CR	Coding Rate
1780	CRC	Cyclic Redundancy Check
1781	DR	Data Rate
1782	ECB	Electronic Code Book
1783	ETSI	European Telecommunications Standards Institute
1784	EIRP	Equivalent Isotropically Radiated Power
1785	FSK	Frequency Shift Keying modulation technique
1786	GPRS	General Packet Radio Service
1787	HAL	Hardware Abstraction Layer
1788	IP	Internet Protocol
1789	LBT	Listen Before Talk
1790	LoRa™	Long Range modulation technique
1791	LoRaWAN™	Long Range Network protocol
1792	MAC	Medium Access Control
1793	MIC	Message Integrity Code
1794	RF	Radio Frequency
1795	RFU	Reserved for Future Usage
1796	Rx	Receiver
1797	RSSI	Received Signal Strength Indicator
1798	SF	Spreading Factor
1799	SNR	Signal Noise Ratio
1800	SPI	Serial Peripheral Interface
1801	SSL	Secure Socket Layer
1802	Tx	Transmitter
1803	USB	Universal Serial Bus
1804		

1805 23 Bibliography**1806 23.1 References**

- 1807 [IEEE802154]: IEEE Standard for Local and Metropolitan Area Networks—Part 15.4: Low-
1808 Rate Wireless Personal Area Networks (LR-WPANs), IEEE Std 802.15.4TM-2011 (Revision
1809 of IEEE Std 802.15.4-2006), September 2011.
- 1810 [RFC4493]: The AES-CMAC Algorithm, June 2006.
- 1811 [RPD1] : LoRa Alliance 2018, RPD_Remote_Multicast_Setup_v1.0.0

1812 24 NOTICE OF USE AND DISCLOSURE

1813 Copyright © LoRa Alliance, Inc. (2017). All Rights Reserved.

1814 The information within this document is the property of the LoRa Alliance (“The Alliance”) and
1815 its use and disclosure are subject to LoRa Alliance Corporate Bylaws, Intellectual Property
1816 Rights (IPR) Policy and Membership Agreements.

1817 Elements of LoRa Alliance specifications may be subject to third party intellectual property
1818 rights, including without limitation, patent, copyright or trademark rights (such a third party may
1819 or may not be a member of LoRa Alliance). The Alliance is not responsible and shall not be
1820 held responsible in any manner for identifying or failing to identify any or all such third party
1821 intellectual property rights.

1822 This document and the information contained herein are provided on an “AS IS” basis and
1823 THE ALLIANCE DISCLAIMS ALL WARRANTIES EXPRESS OR IMPLIED, INCLUDING BUT
1824 NOT LIMITED TO (A) ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN
1825 WILL NOT INFRINGE ANY RIGHTS OF THIRD PARTIES (INCLUDING WITHOUT
1826 LIMITATION ANY INTELLECTUAL PROPERTY RIGHTS INCLUDING PATENT,
1827 COPYRIGHT OR TRADEMARK RIGHTS) OR (B) ANY IMPLIED WARRANTIES OF
1828 MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE OR
1829 NONINFRINGEMENT.

1830 IN NO EVENT WILL THE ALLIANCE BE LIABLE FOR ANY LOSS OF PROFITS, LOSS OF
1831 BUSINESS, LOSS OF USE OF DATA, INTERRUPTION OF BUSINESS, OR FOR ANY
1832 OTHER DIRECT, INDIRECT, SPECIAL OR EXEMPLARY, INCIDENTAL, PUNITIVE OR
1833 CONSEQUENTIAL DAMAGES OF ANY KIND, IN CONTRACT OR IN TORT, IN
1834 CONNECTION WITH THIS DOCUMENT OR THE INFORMATION CONTAINED HEREIN,
1835 EVEN IF ADVISED OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

1836 The above notice and this paragraph must be included on all copies of this document that are
1837 made.

1838 LoRa Alliance, Inc.
1839 3855 SW 153rd Drive
1840 Beaverton, OR 97007

1841 Note: All Company, brand and product names may be trademarks that are the sole property
1842 of their respective owners.

1843