



An In-depth Look at LoRaWAN[®] Class B Devices

Semtech Corporation

November 2019

Introduction

In a LoRaWAN® network, end devices operate in one of three modes: LoRaWAN Class A, Class B, and Class C. As described in *An in-depth look at LoRaWAN Class A devices*, the network can only send messages (*downlinks*) to an end device in Class A mode during one of two short receive windows, which open just after the device has sent a message (an *uplink*) to the network. These uplinks, however, are not prescheduled and may be transmitted by the device at unpredictable times. While this is great for conserving power, that power conservation comes at a cost. For example, end devices in Class A mode do not allow for a known reaction time when the customer application or the server wants to send a message to the device.

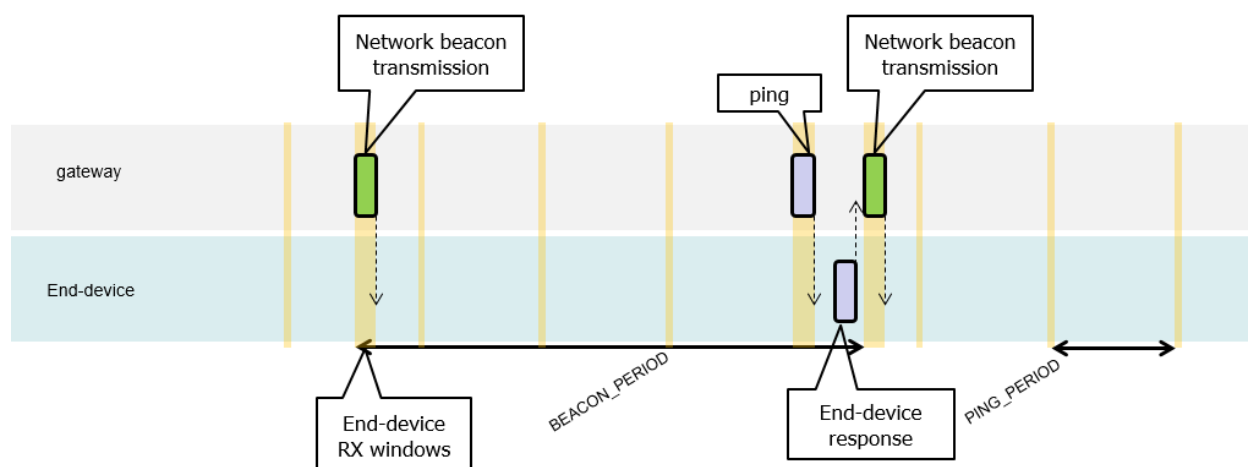
This is where LoRaWAN Class B mode can help. An enhancement of Class A, Class B mode offers regularly-scheduled, fixed-time opportunities for an end device to receive downlinks from the network. All LoRaWAN end devices start in Class A mode; however, devices programmed with a Class B stack during manufacturing may be switched to Class B mode by the application layer.

Class B: The “Beaconing” Class

End devices in Class B mode provide for regularly-scheduled receive windows, in addition to those that open whenever a Class A-style uplink is sent to the server. For this to work, a time-synchronized beacon is broadcast periodically by the network via the gateways, as illustrated in Figure 1. The end device must periodically receive one of these network beacons so that it can align its internal clock with the network.

Based on the beacon timing reference, end devices can open receive windows (*ping slots*) periodically. Any of these ping slots may be used by the network infrastructure to initiate a downlink communication.

Figure 1: Network beacon opportunities

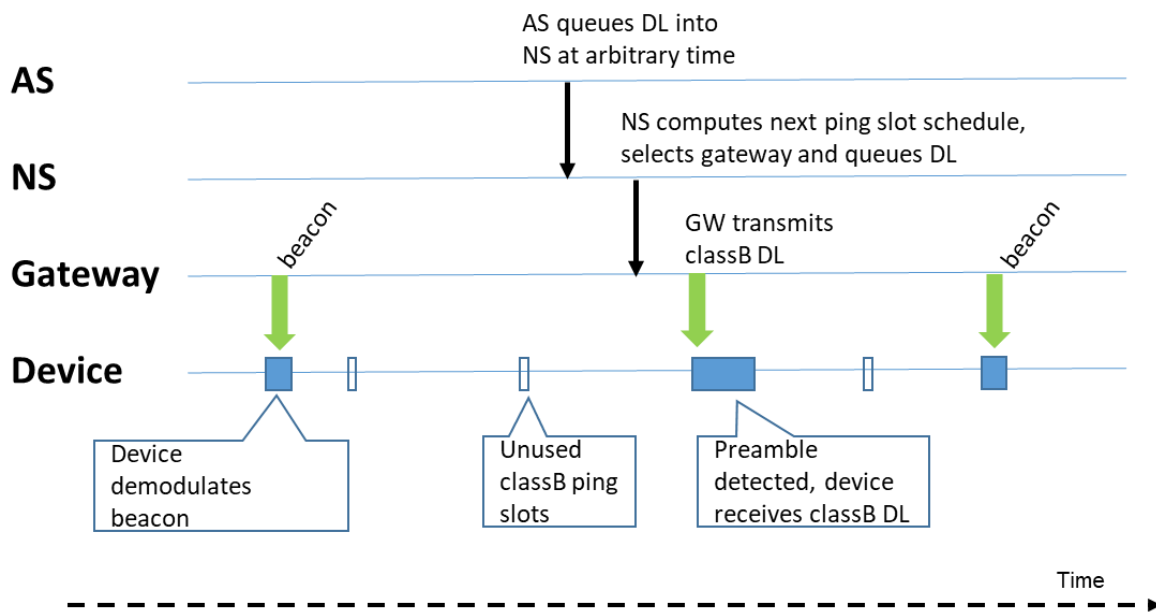


The transmission of a downlink message from an application server to a Class B device is illustrated in Figure 2. Here is a summary of the flow:

1. The application server (AS) queues a downlink message (DL) into the network server (NS)
2. The network server computes the next ping slot schedule
3. The network server computes the best gateway to use based on the last uplink received from the device and the current gateway's transmission schedule
4. The network server queues the downlink into the selected gateway
5. When the selected ping slot start time is reached, the gateway transmits the downlink
6. At the same time, the device turns its receiver on and receives the downlink

When there is no downlink to be sent (roughly 99 percent of the time), the network does not transmit anything. However, the device still opens its receiver at the beginning of the ping slot. When it does not detect a preamble, it goes back to sleep as quickly as possible to conserve power.

Figure 2: Beacon flow



The reception of the synchronization beacons and the ping slots add additional power consumption overhead compared to a Class A device; therefore, a device is allowed to toggle between the Class A (the lowest power consumption mode) and Class B modes.

The decision to switch between the two modes is use-case specific and is left to the device's application layer.

If the switch needs to be controlled from the network side, the customer application must use one of the Class A uplinks to trigger a downlink to the application layer. The application layer on the end device must then recognize and respond to the request to switch modes—this process is not managed at the LoRaWAN level.

For a LoRaWAN network to support Class B end devices, gateways must be able to synchronously broadcast a beacon that provides a timing reference for the end devices' internal clocks. The end devices then use this timing reference to schedule when ping slots are opened to receive potential downlinks from the network.

Because all gateways broadcast the Class B beacon synchronously on the same channel and using the same radio parameters, a device within range of several gateways will receive the superposition of multiple beacons, all experiencing different attenuations and phase distortions.

In order for this end device to be able to demodulate the superposed beacons with as little interference as possible, the gateways must synchronously transmit their beacons with a timing jitter smaller than one microsecond.


This way, the various beacons that superpose at the antenna of the device actually look like a single beacon packet suffering from radio multi-path, therefore they can be demodulated.

The Beaconsing Period

Synchronization beacons are transmitted by network gateways once every 128 seconds.

The interval between two consecutive beacons is divided evenly into 4096 ping slots.

Figure 3: Ping slots



The 128-second beaconsing period has been chosen as a trade-off between minimizing the gateway transmit duty-cycle (saving the gateway transmit time budget for useful downlinks from the application server), while minimizing the power consumption of the end device as it attempts to lock-on to and track the beacon. (To acquire the beacon, the device must keep its receiver on for at least one beacon period.)

You can estimate the power consumption of a Class B end device by first calculating the time on air as well as the time the receive windows are active.

Class B Power Estimation

Here is an example of how to estimate the power consumption for Class B end devices.

The power overhead of Class B operation is a function of:

- The time-on-air of the beacon, which is region-specific
- The periodicity of the ping slots (application driven)
- The ping slot data rate (which sets the minimum duration of each ping slot)
- The average Class B downlink periodicity

For this numerical example, we will assume the following hypothesis:

- Beacon is transmitted using SF9/125kHz => beacon time-on-air ~ 160mSec
- Ping slot periodicity: 32sec (the device opens a ping slot every 16 seconds)
- Default ping slot data rate: SF9/125kHz
- Device receives, on average, one 30-byte Class B downlink per hour.
- The device current in Rx mode is 10mA (typical for SX1272/76 products)

Equation 1: Estimating Class B device power consumption

Beacon TOA	0.16	sec
Beacon periodicity	128	sec
Average beacon demodulation current	12.5	uA
ping slot min duration	32	mSec
ping slot periodicity	32	sec
Average ping slot Rx current	10	uA
ping slot duration when receiving downlink	300	mSec
downlink periodicity	3600	sec
Average Class B downlink demodulation current	0.83	uA
Class B related current	23.33	uA

We can see from this example that for a ping slot periodicity of 32 sec (meaning that the device is reachable with an average latency of 16sec), the power contribution is, roughly, split equally between the beacon synchronization and the opening of the ping slots.

The table above takes into account a possible imprecision of +/-30ppm RTC XTAL. 30ppm over 128sec leads to a +/- 4mSec maximum timing inaccuracy. The 32mSec ping slot duration is long enough to allow a +/- 4mSec misalignment between the transmitter (network) and the receiver (device). For more information about this calculation, see [Application Note AN1200.23](#).

What Does the Beacon Look Like?

All beacons are transmitted using the LoRa modulation in a mode known as *implicit header* mode; that is, without a LoRa® physical header, and without CRC appended by the radio. This is possible because the beacon uses a fixed coding rate and a fixed payload length. The beacon therefore only consists of a preamble and a fixed-length payload.

The beacon preamble is longer than the default. This allows end-devices to implement a low-power, duty-cycled beacon search. This helps to minimize the power consumption of the device during beacon acquisition.

The beacon payload, *BCNPayload*, consists of three parts: a network common part (that is, this data is sent identically by all gateways), an optional gateway-specific part (this part may be different for each gateway), and a timestamp time in seconds since January 6, 1980 00:00:00 UTC (the start of the GPS epoch). The integrity of the beacon's network common part is protected by a 16-bit CRC. For specific information about how the CRC is computed, see the latest [LoRaWAN specification](#).

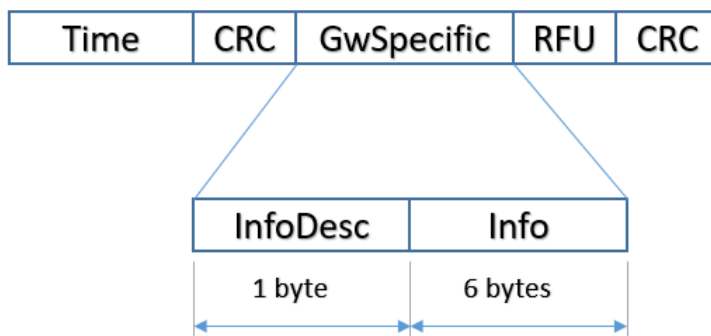
The RFU field is equal to zero. (The length of the RFU field is region-specific.)

Note: In practice, many networks only broadcast the network common part of the beacon and do not include gateway-specific fields.

The Gateway-Specific Beacon Element

The gateway-specific portion of the beacon payload, when present, provides additional information regarding the gateway that is sending the beacon. Therefore, this portion of the beacon will likely be different for each gateway. The gateway-specific portion of the beacon payload is protected by a CRC-16 computed on the *GwSpecific+RFU* fields. The CRC-16 definition is the same as that of the network common part.

Figure 4: Beacon payload



The InfoDesc and Info Fields

The information descriptor, *InfoDesc*, describes how the information field should be interpreted. For a single omnidirectional antenna gateway, the *InfoDesc* value is 0 when broadcasting GPS coordinates. In contrast, for a site featuring sectored antennas for example, the first antenna broadcasts the beacon with *InfoDesc* equal to 0, and the second antenna with *InfoDesc* field equal to 1, etc. If *InfoDesc* equals 0, 1 or 2, the content of the *Info* field encodes the GPS coordinates of the antenna broadcasting the beacon.


Table 1: InfoDesc field values

InfoDesc	Info
----------	------

Value	Description
0	GPS coordinate of the gateway's first antenna
1	GPS coordinate of the gateway's second antenna
2	GPS coordinate of the gateway's third antenna
3:127	RFU
128:256	Reserved for custom, network-specific broadcasts

The latitude and longitude (*Lat* and *Lng*) fields within the *Info* field encode the geographical location of the gateway.

Figure 5: Info field format



The North-South latitude is encoded using a signed, 24-bit word, where -2^{23} corresponds to 90° south (the South Pole) and 2^{23} corresponds to 90° north (the North Pole). The equator corresponds to 0.

The East-West longitude is encoded using a signed 24-bit word where -2^{23} corresponds to 180° west and 2^{23} corresponds to 180° east. The Greenwich meridian corresponds to 0.

Beacon Acquisition and Tracking

As mentioned previously, an end device must receive a timing beacon to bring its internal time base in sync with the network before it can switch from Class A mode to Class B. Once successfully operating in Class B mode, the end device must stay in alignment with the network beacons. This helps ensure that the device's internal time base does not drift compared to the network time.

A device operating in Class B mode may temporarily stop receiving beacons for many reasons (interference, device moving in and out of the network coverage area, etc.).

When the device temporarily loses a beacon, it relies on its internal clock to keep opening the Class B synchronous ping slots. The LoRaWAN specification requires that Class B devices should be able to operate in Class B mode without receiving the beacon for up to two hours.

This temporary Class B operation without a beacon is called *beaconless* operation. During this holding period, the reception of any beacon directed to the end device will extend beaconless operation for another two hours.

If a Class B end device is temporarily unable to receive beacons, it will gradually widen its beacon and ping slot reception windows to account for any drift of its internal time base. Obviously, a higher precision time base allows the device to open narrower reception windows, thus minimizing its power consumption. The accuracy of the real time clock oscillator is therefore very important for Class B devices. If a Class B end device doesn't receive a beacon during the 120-minute holdover period, the end device will return to operating in Class A mode.


Ping Slots and Ping Periods

Based on the beacon timing reference, end devices periodically open *ping slots* (reception windows). A network-initiated downlink to an end device using one of these ping slots is called a *ping*. The gateway chosen to transmit this ping is selected by the network server based on signal quality indicators from the last uplink sent by the end device. For this reason, mobile Class B devices must send periodic uplinks, so that the network server can update the downlink routing-path database.

To avoid systematic collisions and problems of overhearing messages, each device's ping slot index is randomized and changed at every beacon period.

The ping slot periodicity may be as short as one second, or as long as 128 seconds, as illustrated in Figure 6.

Figure 6: Ping slot periodicity



The device opens ping slots regularly. However, the network does not always have downlinks to transmit. Therefore, most ping slots are not used by the network. When this is the case, the reception window on the end device is closed as soon as the radio transceiver determines that no preamble is present. If a preamble *is* detected, the radio transceiver will remain on until the downlink frame is demodulated. The MAC layer will then process the frame, check that the address field matches the end device address, and that the Message Integrity Check (MIC) is valid. Only then will the message be forwarded to the application layer on the end device.

Unicast and Multicast Pings

Downlinks can be either *unicast* or *multicast*. Unicast messages are sent to a single end device, while multicast messages are sent to several end devices at the same time. For a multicast message to be successful, the following conditions must be met:

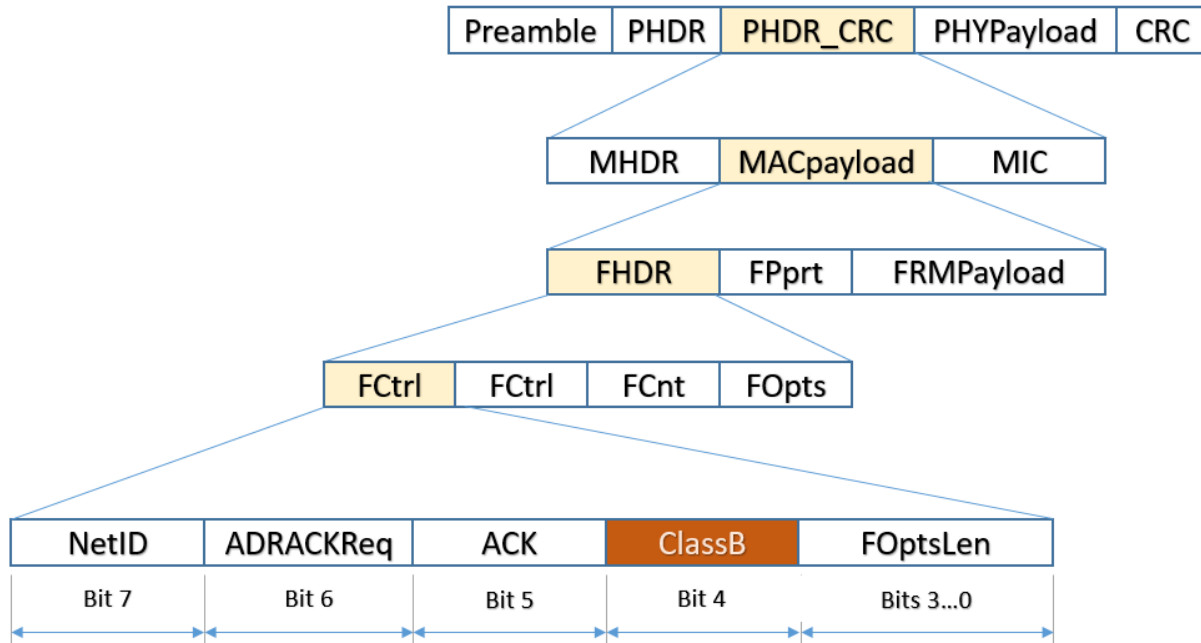
- All devices in a multicast group must share the same multicast address and associated encryption keys
- All devices must be open at the same time, on the same channel, using the same data rate

Note: To remotely set up a multicast group please see the [LoRaWAN Remote Multicast Setup Specification \(v1.0.0\)](#)

Class B: Uplink Frame and MAC Commands

Uplink frames in Class B mode have the same structure as those used in Class A communications, with the exception of the Frame Control (*FCtrl*) octet in the frame header.

Figure 7: Uplink frame structure



For Class A communication uplinks, Bit 4 of the *FCtrl* octet is set to 0. This bit is set to 1 for Class B uplinks. Setting the Class B bit to 1 in an uplink message will signal to the network server that the device has switched to Class B mode and is ready to receive scheduled downlink pings.

Because a LoRa-based device always starts in Class A mode, All MAC commands specified for Class A end devices MUST be implemented in Class B end devices, in addition to those commands specifically used for Class B operation. For more information, see the [LoRaWAN Specification, v1.0.3](#).

Conclusion

By design, all LoRa-based end devices support Class A communication. Class A devices can send uplinks to the network server at any time, and only listen for downlinks immediately after sending their uplinks. In contrast, Class B devices allow for predictable times when the network server can send a downlink to the device. The trade-off for this predictability is that Class B devices are not as energy efficient as those in Class A mode; therefore, the battery life of Class B devices is shorter than that of Class A devices.



Important Notice

Information relating to this product and the application or design described herein is believed to be reliable, however such information is provided as a guide only and Semtech assumes no liability for any errors in this document, or for the application or design described herein. Semtech reserves the right to make changes to the product or this document at any time without notice. Buyers should obtain the latest relevant information before placing order and should verify that such information is current and complete. Semtech warrants performance of its products to the specifications applicable at the time of sale, and all sales are made in accordance with Semtech's standard terms and conditions of sale.

SEMTECH PRODUCTS ARE NOT DESIGNED, INTENDED, AUTHORIZED OR WARRANTED TO BE SUITABLE FOR USE IN LIFE-SUPPORT APPLICATIONS, DEVICES OR SYSTEMS, OR IN NUCLEAR APPLICATIONS IN WHICH THE FAILURE COULD BE REASONABLY EXPECTED TO RESULT IN PERSONAL INJURY, LOSS OF LIFE OR SEVERE PROPERTY OR ENVIRONMENTAL DAMAGE. INCLUSION OF SEMTECH PRODUCTS IN SUCH APPLICATIONS IS UNDERSTOOD TO BE UNDERTAKEN SOLELY AT THE CUSTOMER'S OWN RISK. Should a customer purchase or use Semtech products for any such unauthorized application, the consumer shall indemnify and hold Semtech and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs, damages and attorney fees which could arise.

The Semtech name and logo are registered trademarks of the Semtech Corporation. All other trademarks and trade names mentioned may be marks and names of Semtech or their respective companies. Semtech reserves the right to make changes to, or discontinue any products described in this document without further notice. Semtech makes no warranty, representation guarantee, express or implied, regarding the suitability of its products for any particular purpose. All rights reserved.

©Semtech 2019

Contact Information

Semtech Corporation
200 Flynn Road, Camarillo, CA 93012
Phone: (805) 498-2111, Fax: (805) 498-3804
www.semtech.com