



LoRa® and LoRaWAN®

Table of Contents

List of figures	2
List of tables.	3
1. What are LoRa® and LoRaWAN®?	4
2. Radio Modulation and LoRa.	7
2.1. Key LoRa Modulation Properties	9
2.2. LoRa Modulation Characteristics	10
2.3. Data Collisions and Spreading Factor Orthogonality	11
3. LoRaWAN Network Fundamentals	13
3.1. LoRaWAN Network Elements: An Introduction	13
3.1.1. LoRa-based End Devices	14
3.1.2. LoRaWAN Gateways	15
3.1.3. Network Server	16
3.1.4. Application Servers	17
3.1.5. Join Server	17
3.2. LoRaWAN Network Elements: Device Commissioning	18
3.3. LoRaWAN Network Elements: Security.	19
3.3.1. The Join Procedure	19
3.4. Device Classes: A, B and C	21
3.4.1. Class A Devices	22
3.4.2. Class B Devices	23
3.4.3. Class C Devices	25
4. The LoRa Alliance	27
Disclaimer	28

List of figures

- Figure 1. IoT Technologies
- Figure 2. Advantages of deploying a LoRaWAN network
- Figure 3. OSI seven-layer network model
- Figure 4. DSSS system carrier phase transmitter signal changes
- Figure 5. LoRa Chirp Spread Spectrum illustration
- Figure 6. LoRa Spreading Factors
- Figure 7. LoRaWAN North America Channel Plan
- Figure 8. LoRa modulation characteristics
- Figure 9. LoRaWAN technology stack
- Figure 10. Typical LoRaWAN network implementation
- Figure 11. End devices in a typical LoRaWAN network deployment
- Figure 12. Gateways in a typical LoRaWAN network deployment
- Figure 13. Gateways receiving and transmitting messages from end devices
- Figure 14. LoRaWAN Network Server in a typical LoRaWAN network deployment
- Figure 15. LoRaWAN Application Server in a typical LoRaWAN network deployment
- Figure 16. LoRaWAN Join Server in a typical LoRaWAN network deployment
- Figure 17. Activation Types
- Figure 18. Security keys generated during the Join procedure
- Figure 19. Sending a join request message to the join server
- Figure 20. Sending a join accept message to an end device
- Figure 21. Session keys are shared with the network server and the application server
- Figure 22. Secure transmission of data packets
- Figure 23. Class A operation
- Figure 24. Class A operation when nothing is received
- Figure 25. Class A operation when a data packet is received in the first receive window
- Figure 26. Class A operation when a data packet is received in the second receive window
- Figure 27. Class B beaconing operations
- Figure 28. Periodic Class B beaconing for device synchronization
- Figure 29. Class B ping slots
- Figure 30. Class C operation

List of tables

Chapter 1. What are LoRa® and LoRaWAN®?

LoRa is an RF modulation technology for low-power, wide area networks (LPWANs). The name, LoRa, is a reference to the extremely long-range data links that this technology enables. Created by Semtech to standardize LPWANs, LoRa provides for long-range communications: up to three miles (five kilometers) in urban areas, and up to 10 miles (15 kilometers) or more in rural areas (line of sight). A key characteristic of the LoRa-based solutions is ultra-low power requirements, which allows for the creation of battery-operated devices that can last for up to 10 years. Deployed in a star topology, a network based on the open LoRaWAN protocol is perfect for applications that require long-range or deep in-building communication among a large number of devices that have low power requirements and that collect small amounts of data.

Consider the differences between LoRa and other network technologies that are typically used in IoT or traditional machine-to-machine (M2M) connectivity solutions:


<u>Traditional Cellular</u> Long Range High Data Rates Low Battery Life High Cost	<div> <div>LPWAN (3-5B in 2022)</div> <div>  <div> Long Range Low Data Rates Long Battery Life Low Cost High Capacity Potential </div> </div> </div>	<u>Cat-M1</u> Long Range High Data Rates Low Battery Life Medium Cost
<u>Local Area Network</u> (Wi-Fi) Short Range High Data Rates Low Battery Life Medium Cost	<u>Narrow-Band IoT</u> (NB-IoT) Stationary Devices Short Range (indoor coverage) Low Data Rates Good Battery Life Low Cost	<u>Personal Area Network</u> (Bluetooth®) Very Short Range Low data rates Good Battery Life Low Cost

Figure 1. IoT Technologies

NOTE

In Europe, mobile network operators have implemented a dual strategy to address packet size and latency issues. They often offer both LoRaWAN and Cat-M1, which are complementary technologies. LoRaWAN accommodates the need for longer battery life, with a trade-off of longer latency and smaller packet sizes. In contrast Cat-M1 can be used for larger payloads with less latency than LoRaWAN can accommodate.

Figure 2 highlights some important advantages of deploying a LoRaWAN network:

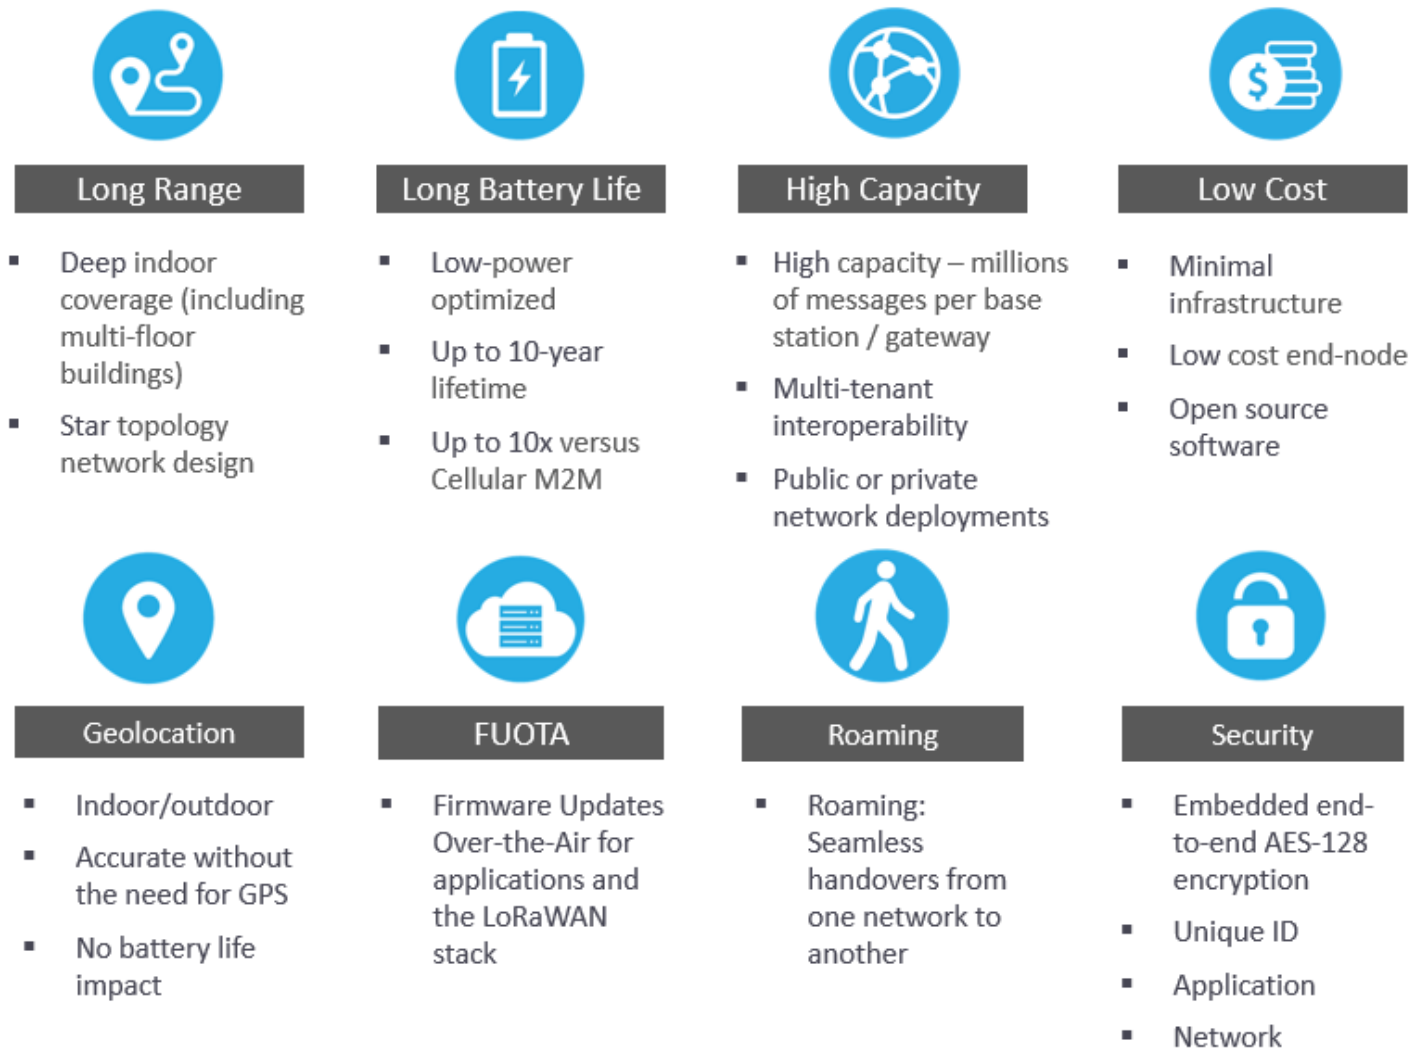


Figure 2. Advantages of deploying a LoRaWAN network

Let's look into these advantages in a little more depth.

With respect to range, a single LoRa-based gateway can receive and transmit signals over a distance of more than 10 miles (15 kilometers) in rural areas. Even in dense urban environments, messages are able to travel up to three miles (five kilometers), depending on how deep indoors the end devices (end nodes) are located.

As far as battery life goes, the energy required to transmit a data packet is quite minimal given that the data packets are very small and only transmitted a few times a day. Furthermore, when the end devices are asleep, the power consumption is measured in milliwatts (mW), allowing a device's battery to last for many, many years.

When it comes to capacity, a LoRaWAN network can support millions of messages. However, the number of messages supported in any given deployment depends upon the number of gateways that are installed. A single eight-channel gateway can support a few hundred thousand messages over the course of a 24-hour period. If each end device sends 10 messages a day, such a gateway can support about 10,000 devices.^[1] If the network includes 10 such gateways, the network can support roughly 100,000 devices and one million messages. If more capacity is required, all that is needed is to add additional gateways to the network.

And then, there is cost. Given the capabilities of LoRa-based end nodes and gateways, only a few gateways - configured in a star network - are required

to serve a multitude of end nodes. This means that capital and operational expenses can be kept relatively low. Also, when the cost-effective LoRa RF modules that are embedded in inexpensive end nodes are used in conjunction with the open LoRaWAN standard, the return on investment can be considerable.

[1] There is no one-to-one relationship between LoRa-based devices and gateways in a LoRaWAN network; messages sent to and from end devices travel through all gateways within range. De-duplication is handled by the network server.

Chapter 2. Radio Modulation and LoRa

A proprietary spread-spectrum modulation technique derived from existing Chirp Spread Spectrum (CSS) technology, LoRa offers a trade-off between sensitivity and data rate, while operating in a fixed-bandwidth channel of either 125 KHz or 500 KHz (for uplink channels), and 500 KHz (for downlink channels). Additionally, LoRa uses orthogonal spreading factors. This allows the network to preserve the battery life of connected end nodes by making adaptive optimizations of an individual end node's power levels and data rates. For example, an end device located close to a gateway should transmit data at a low spreading factor, since very little link budget is needed. However, an end device located several miles from a gateway will need to transmit with a much higher spreading factor. This higher spreading factor provides increased processing gain, and higher reception sensitivity, although the data rate will, necessarily, be lower.

LoRa is purely a physical (PHY), or “bits” layer implementation, as defined by the OSI seven-layer Network Model, depicted in [Figure 3](#). Instead of cabling, the air is used as a medium for transporting LoRa radio waves from an RF transmitter in an IoT device to an RF receiver in a gateway, and vice versa.

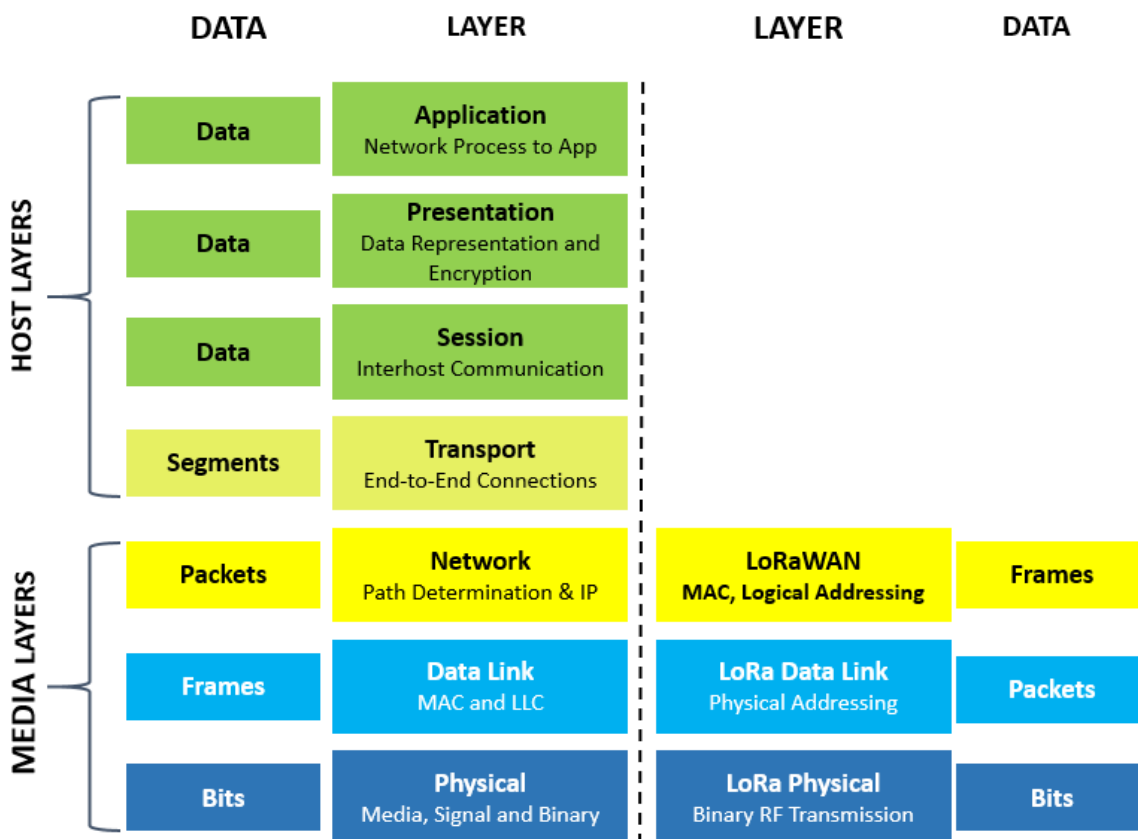


Figure 3. OSI seven-layer network model

In a traditional or Direct Sequence Spread Spectrum (DSSS) system, the carrier phase of the transmitter signal changes according to a code sequence as shown in [Figure 4](#). When multiplying the data signal with a pre-defined bit pattern at a much higher rate, also known as a spreading code (or chip sequence), a “faster” signal is created that has higher frequency components than the original data signal. This means that the signal bandwidth is spread beyond the bandwidth of the original signal. In RF terminology, the bits of the code sequence are called chips (in order to distinguish between the longer, un-coded, bits of the original data signal). When the transmitted signal arrives at the RF receiver, it is multiplied with an identical copy of the spreading code used in the RF transmitter, resulting in a replica of the original data signal.

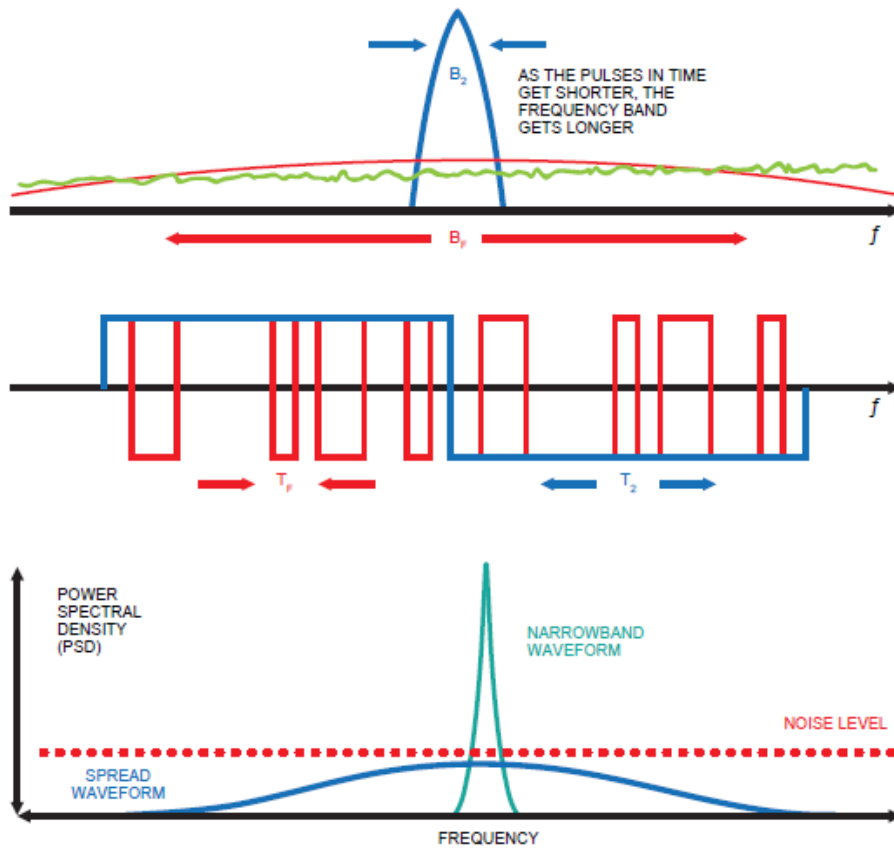


Figure 4. DSSS system carrier phase transmitter signal changes

You might ask: Why go through all this trouble? Why not just transmit the original data signal instead of going through this code sequence multiplication? The answer is simple: going through this code sequence multiplication buys you a higher RF link budget, so you can transmit over a longer range.

The Log10 ratio of the code sequence's chip rate and the data signal's bit rate is called the processing gain (G_p). This gain is what allows the receiver to recover the original data signal, even if the channel has a negative signal-to-noise ratio (SNR). LoRa has a superior G_p compared to frequency-shift keying (FSK) modulation, allowing for a reduced transmitter output power level while maintaining the same signal data rate and a similar link budget.

One of the downsides of a DSSS system is the fact that it requires a highly-accurate (and expensive) reference clock. Semtech's LoRa Chirp Spread Spectrum (CSS) technology offers a low-cost and low-power, yet robust, DSSS alternative that does not require a highly-accurate reference clock. In LoRa modulation, the spreading of the signal's spectrum is achieved by generating a chirp signal that continuously varies in frequency, as is depicted in [Figure 5](#).




Figure 5. LoRa Chirp Spread Spectrum illustration

An advantage of this method is that the timing and frequency offsets between transmitter and receiver are equivalent, greatly reducing the complexity of the receiver design. The frequency bandwidth of this chirp is equivalent to the spectral bandwidth of the signal. The data signal that carries the data from an end device to a gateway is chipped at a higher data rate and modulated onto the chirp carrier signal. LoRa modulation also includes a variable error correction scheme that improves the robustness of the transmitted signal. For every four bits of information sent, a fifth bit of parity information is sent.

2.1. Key LoRa Modulation Properties

As noted above, LoRa processing gain is introduced in the RF channel by multiplying the data signal with a spreading code or chip sequence. By increasing the chip rate, we increase the frequency components of the total signal spectrum. In other words, the energy of the total signal is now spread over a wider range of frequencies, allowing the receiver to discern a signal with a lower (that is, worse) signal-to-noise ratio (SNR).

In LoRa terms, the amount of spreading code applied to the original data signal is called the *spreading factor* (SF). LoRa modulation has a total of six spreading factors (SF7 to SF12). The larger the spreading factor used, the farther the signal will be able to travel and still be received without errors by the RF receiver.

Figure 6 shows the four different spreading factors [SF7...SF10] that can be used for uplink (UL) messages on a 125 KHz channel.^[4] It shows the equivalent bit rate as well as the estimated range (this depends on the terrain; longer distances will be achieved in a rural environment than in an urban environment). It also shows the dwell time, or *time on air* (TOA), values for an 11-byte payload for each of the four spreading factors.

Spreading Factor (For UL at 125 KHz)	Bit Rate	Range (Depends on Terrain)	Time on Air for an 11-byte payload
SF10	980 bps	8 km	371 ms
SF9	1760 bps	6 km	185 ms
SF8	3125 bps	4 km	103 ms
SF7	5470 bps	2 km	61 ms

Figure 6. LoRa Spreading Factors

Importantly, the LoRa modulation spreading factors are inherently orthogonal. This means that signals modulated with different spreading factors and transmitted on the same frequency channel at the same time do not interfere with each other. Instead, signals at different spreading factors simply appear to be noise to each other.

LoRa signals are robust and very resistant to both in-band and out-of-band interference mechanisms. LoRa modulation also offers immunity to multipath and fading, making it ideal for use in urban and suburban environments, where both mechanisms dominate. Additionally, Doppler shifts cause a small frequency shift in the time axis of the baseband signal. This frequency offset tolerance mitigates the requirement for tight tolerance reference clock sources and, therefore, makes LoRa ideal for data communications from devices that are mobile.

2.2. LoRa Modulation Characteristics

The LoRa modulation characteristics for each region are defined in the LoRaWAN Regional Parameters document, available from the LoRa Alliance®. In North America, there are 64, 125 kHz LoRa uplink channels defined, centered on a 200 kHz raster as can be seen in [Figure 7](#). There are eight 500 kHz uplink channels as well as eight, 500 kHz downlink channels defined. In North America, gateways can have up to 64, 125 kHz uplink channels as well as eight 500 kHz uplink and downlink channels. This type of gateway is referred to as a carrier grade macro gateway and is used for outdoor applications only.




Figure 7. LoRaWAN North America Channel Plan

[Figure 8](#) provides another way to understand these modulation characteristics.

Data Rate (DR)	Spreading Factor (SF)	Channel Frequency	Uplink or Downlink	Bitrate (Bits/Sec)	Maximum User Payload Size (Bytes)
0	SF10	125 kHz	Uplink	980	11
1	SF9	125 kHz	Uplink	1,760	53
2	SF8	125 kHz	Uplink	3,125	125
3	SF7	125 kHz	Uplink	5,470	242
4	SF8	500 kHz	Uplink	12,500	242
5 – 7					
8	SF12	500 kHz	Downlink	980	53
9	SF11	500 kHz	Downlink	1,760	129
10	SF10	500 kHz	Downlink	3,125	242
11	SF9	500 kHz	Downlink	5,470	242
12	SF8	500 kHz	Downlink	12,500	242
13	SF8	500 kHz	Downlink	21,900	242

Figure 8. LoRa modulation characteristics

- The LoRa physical layer is intended for low throughput, low data rate, and high link budget (i.e., “long-range”) applications.
- For a fixed channel bandwidth, the higher the spreading factor, the higher the processing gain, resulting in an increase in sensitivity and, therefore, an increase in link budget. Subsequently, however, the time on air will also increase.
- Orthogonality between spreading factors allows for the transmission of multiple LoRa signals that are both on the same channel frequency **and** in the same time-slot.
- For a fixed SF, a narrower bandwidth will increase sensitivity as the bit rate is reduced.
- LoRaWAN in North America uses 125 kHz uplink channels and 500 kHz uplink and downlink channels
- The Code Rate is the degree of redundancy implemented by the forward error correction (FEC) used to detect errors and correct them. This rate is fixed at 4/5 for the LoRaWAN protocol

As Stephan Hengstler asserts in his book, *A Novel Chirp Modulation Spread Spectrum technique for Multiple Access*, “LoRa is a constant envelope modulation (very low cost, power efficient power amplifier implementation) ... [it] is the most robust, ultra-low power and long range RF solution available.”

2.3. Data Collisions and Spreading Factor Orthogonality

With LoRa, packets using different spreading factors are orthogonal, meaning that they are invisible to each other: as mentioned earlier, they simply appear as noise to one another. Therefore, two packets that arrive at the same time on the same receive channel at different spreading factors will not collide and, both will be demodulated by the gateway modem chip. However, two packets with the *same* spreading factor arriving at the same time on the same channel **might** result in a collision. However if one of the two packets is stronger by six dB, it will survive.

The capacity of a LoRaWAN network is a function of its gateway density. To maximize the capacity of the network, using an adaptive data rate (ADR) mechanism is essential. The main goal of ADR is to save the battery power of the LoRaWAN end-nodes. By having the end-nodes closest to a gateway transmit using the lowest spreading factor, their time on air is minimized, thereby prolonging their battery life. More distant sensors transmit at a higher

spreading factor. A trade-off is made between battery power and distance given that a higher spreading factor allows for a gateway to connect to devices that are farther away.

[1] Downlink messages broadcast over 500 KHz channels can use all six available spreading factors (SF7...SF12).

Chapter 3. LoRaWAN Network Fundamentals

To fully understand LoRaWAN networks, we will start with a look at the technology stack. As shown in [Figure 9](#), LoRa is the physical (PHY) layer, i.e., the wireless modulation used to create the long-range communication link. LoRaWAN is an open networking protocol that delivers secure bi-directional communication, mobility, and localization services standardized and maintained by the LoRa Alliance.

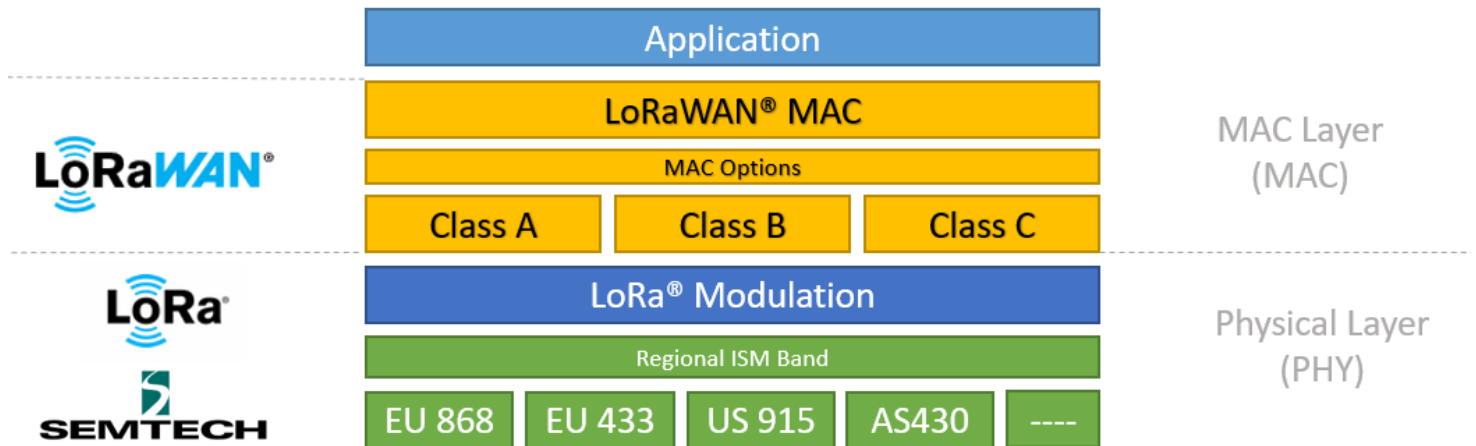


Figure 9. LoRaWAN technology stack

3.1. LoRaWAN Network Elements: An Introduction

Now that we have a basic understanding of LoRa, we will examine the architecture of a LoRaWAN network. [Figure 10](#) shows a typical LoRaWAN network implementation from end to end.

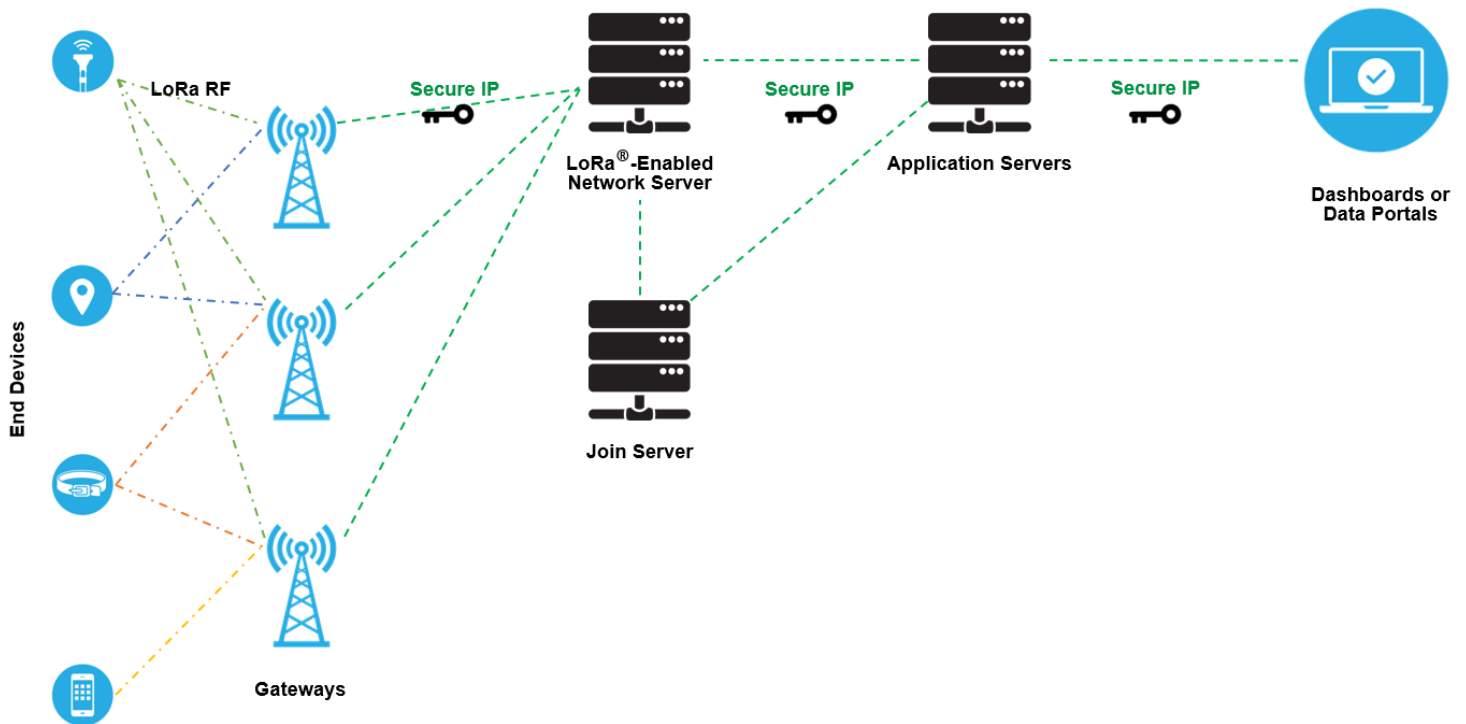


Figure 10. Typical LoRaWAN network implementation

Let us examine this diagram in smaller pieces.

3.1.1. LoRa-based End Devices

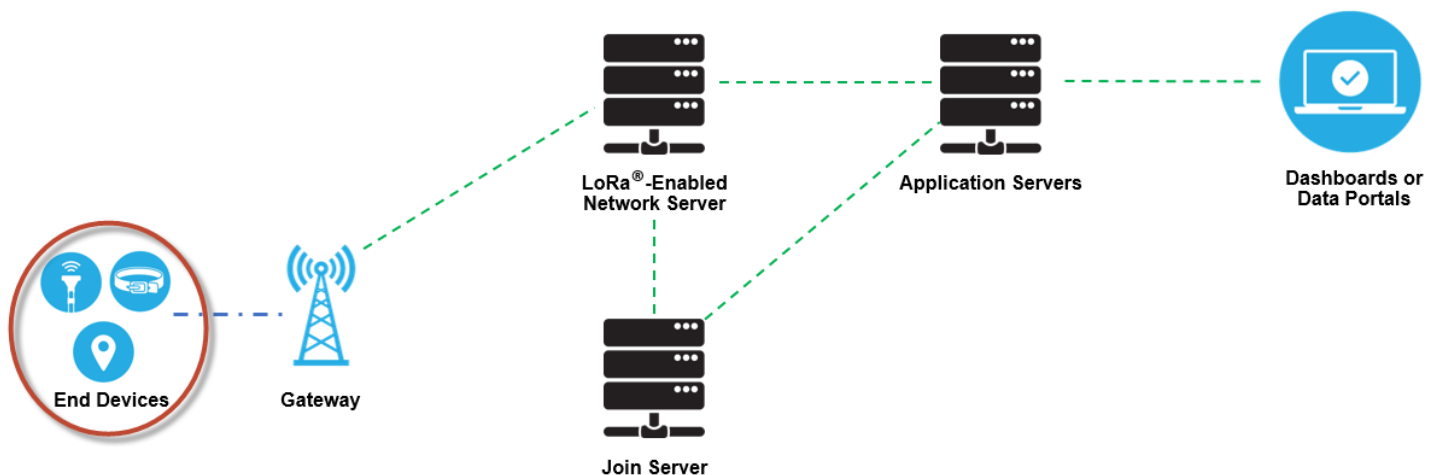


Figure 11. End devices in a typical LoRaWAN network deployment

A LoRaWAN-enabled **end device** is a sensor or an actuator which is wirelessly connected to a LoRaWAN network through radio gateways using LoRa RF Modulation.

In the majority of applications, an end device is an autonomous, often battery-operated sensor that digitizes physical conditions and environmental events. Typical use cases for an actuator include: street lighting, wireless locks, water valve shut off, leak prevention, among others.

When they are being manufactured, LoRa-based devices are assigned several unique identifiers. These identifiers are used to securely activate and administer the device, to ensure the safe transport of packets over a private or public network and to deliver encrypted data to the Cloud.

3.1.2. LoRaWAN Gateways




Figure 12. Gateways in a typical LoRaWAN network deployment

A LoRaWAN **gateway** receives LoRa modulated RF messages from any end device in hearing distance and forwards these data messages to the LoRaWAN network server (LNS), which is connected through an IP backbone. There is no fixed association between an end device and a specific gateway. Instead, the same sensor can be served by multiple gateways in the area. With LoRaWAN, each uplink packet sent by the end-device will be received by all gateways within reach, as illustrated in Figure 12. This arrangement significantly reduces packet error rate (since the chances that at least one gateway will receive the message are very high), significantly reduces battery overhead for mobile/nomadic sensors, and allows for low-cost geolocation (assuming the gateways in question are geolocation-capable).

The IP traffic from a gateway to the network server can be backhauled via Wi-Fi, hardwired Ethernet or via a Cellular connection. LoRaWAN gateways operate entirely at the physical layer and, in essence, are nothing but LoRa radio message forwarders. They only check the data integrity of each incoming LoRa RF message. If the integrity is not intact, that is, if the CRC is incorrect, the message will be dropped. If correct the gateway will forward it to the LNS, together with some metadata that includes the receive RSSI level of the message as well as an optional timestamp. For LoRaWAN downlinks, a gateway executes transmission requests coming from the LNS without any interpretation of the payload. Since multiple gateways can receive the same LoRa RF message from a single end device, the LNS performs data de-duplication and deletes all copies. Based on the RSSI levels of the identical messages, the network server typically selects the gateway that received the message with the best RSSI when transmitting a downlink message because that gateway is the one closest to the end device in question.

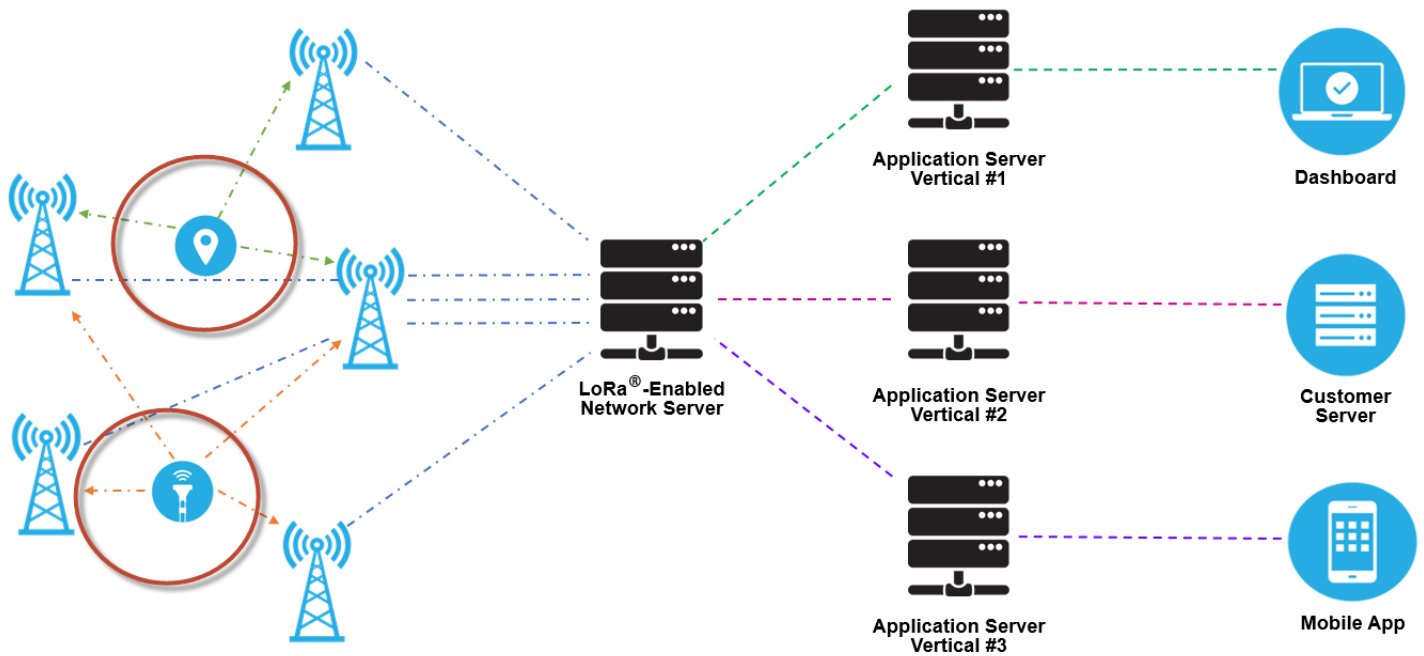


Figure 13. Gateways receiving and transmitting messages from end devices

Furthermore, LoRa allows for scalable, cost-optimized gateway implementation, depending on deployment objectives. For example, in North America, 8-, 16-, and 64-channel gateways are available.

The 8-channel gateways are the least expensive. The type of gateway needed will depend on the use case. Eight- and 16-channel gateways are available for both indoor and outdoor use. Sixty-four channel gateways are only available in a carrier-grade variant. This type of gateway is intended for deployment in such places as cell towers, the rooftops of very tall buildings, etc.

3.1.3. Network Server




Figure 14. LoRaWAN Network Server in a typical LoRaWAN network deployment

The LoRaWAN network server (LNS) manages the entire network, dynamically controls the network parameters to adapt the system to ever-changing conditions, and establishes secure 128-bit AES connections for the transport of both the end to end data (from LoRaWAN end device to the end users

Application in the Cloud) as well as for the control of traffic that flows from the LoRaWAN end device to the LNS (and back). The network server ensures the authenticity of every sensor on the network and the integrity of every message. At the same time, the network server cannot see or access the application data.

In general, all LoRaWAN network servers share the following features:

- Device address checking
- Frame authentication and frame counter management
- Acknowledgements of received messages
- Adapting data rates using the ADR protocol
- Responding to all MAC layer requests coming from the device,
- Forwarding uplink application payloads to the appropriate application servers
- Queuing of downlink payloads coming from any Application Server to any device connected to the network
- Forwarding Join-request and Join-accept messages between the devices and the join server

3.1.4. Application Servers

Application servers are responsible for securely handling, managing and interpreting sensor application data. They also generate all the application-layer downlink payloads to the connected end devices.




Figure 15. LoRaWAN Application Server in a typical LoRaWAN network deployment

3.1.5. Join Server

The join server manages the over-the-air activation process for end devices to be added to the network.

The join server contains the information required to process uplink *join-request* frames and generate the downlink *join-accept* frames. It signals to the network server which application server should be connected to the end-device, and performs the network and application session encryption key derivations. It communicates the Network Session Key of the device to the network server, and the Application Session Key to the corresponding application server.




Figure 16. LoRaWAN Join Server in a typical LoRaWAN network deployment

For that purpose, the join server must contain the following information for each end-device under its control:

- DevEUI (end-device serial unique identifier)
- AppKey (application encryption key)
- NwkKey (network encryption key)
- Application Server identifier
- End-Device Service Profile

3.2. LoRaWAN Network Elements: Device Commissioning

For the sake of security, quality of service, billing, and other needs, devices must be commissioned and activated on the network at the start of operation. The commissioning process securely aligns each device and the network with respect to essential provisioning parameters (such as identifiers, encryption keys, and server locations)

The LoRaWAN specification allows for two types of activation: Over-the-Air Activation (OTAA) (preferred) and Activation by Personalization (ABP). [Figure 17](#) shows the different characteristics of each of these types of activation.

Over-the-Air Activation (OTAA)	Activation by Personalization (ABP)
<ul style="list-style-type: none"> • Device manufacturers autonomously generate essential provisioning parameters • Secure keys (session-long and derived) can be renewed regularly • Devices can store multiple “identities” to dynamically and securely switch networks and operators during its lifetime • High-grade, tamper-proof security options are available 	<ul style="list-style-type: none"> • A simplified (less secure) commissioning process • IDs and Keys are personalized at fabrication • Devices become immediately functional upon powering up; the Join procedure is skipped • Devices are tied to a specific network/service; the NetID is a portion of the device network address

Figure 17. Activation Types

3.3. LoRaWAN Network Elements: Security

There are two key elements to the security of a LoRaWAN network: the *join procedure* and *message authentication*. The join procedure establishes mutual authentication between an end device and the LoRaWAN network to which it is connected. Only authorized devices are allowed to join the network. LoRaWAN MAC and application messages are origin-authenticated, integrity-protected and encrypted end-to-end (i.e., from end device to the application server and vice versa).

These security features ensure that:

- Network traffic has not been altered
- Only legitimate devices are connected to the LoRaWAN network
- Network traffic cannot be listened to (no eavesdropping)
- Network traffic cannot be captured and replayed

With that foundation, we will take a look at the LoRaWAN security measures in more detail.

3.3.1. The Join Procedure

We will begin with the security keys, as illustrated in [Figure 18](#). Individual root keys are securely stored on the end devices, and matching keys are securely stored on the join server.




Figure 18. Security keys generated during the Join procedure

The end device sends a *join request* message to the join server, as illustrated in [Figure 19](#).

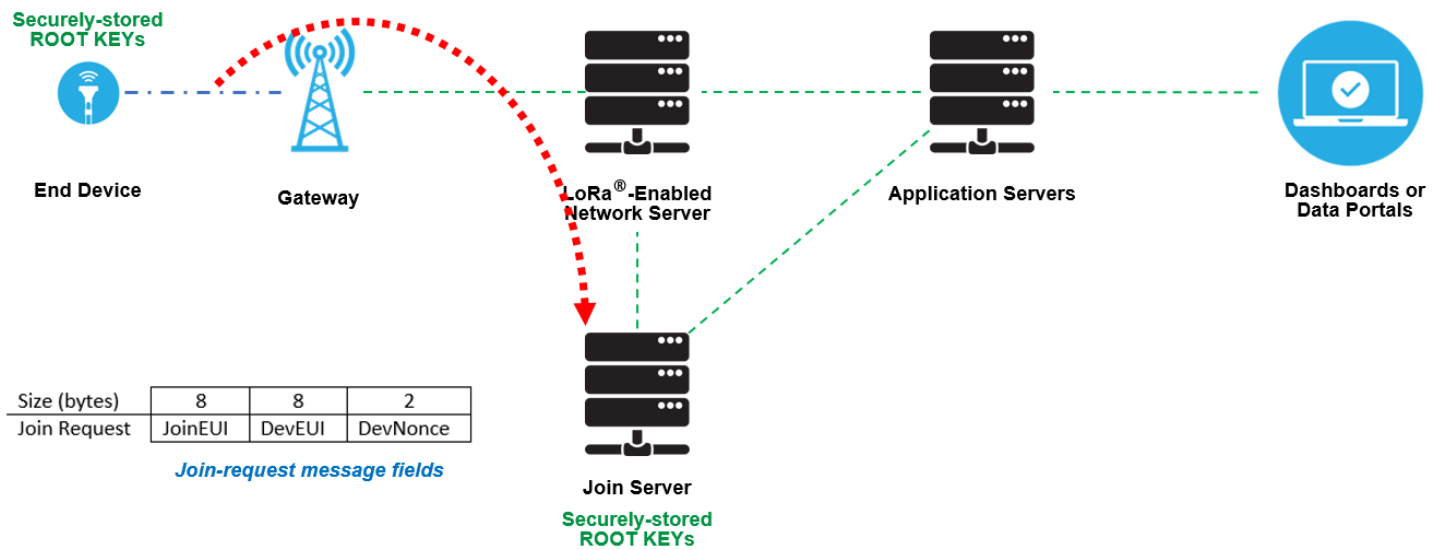


Figure 19. Sending a join request message to the join server

After the join server authenticates the device requesting to join the network, it returns a *join accept* message to the device, as illustrated in Figure 20.

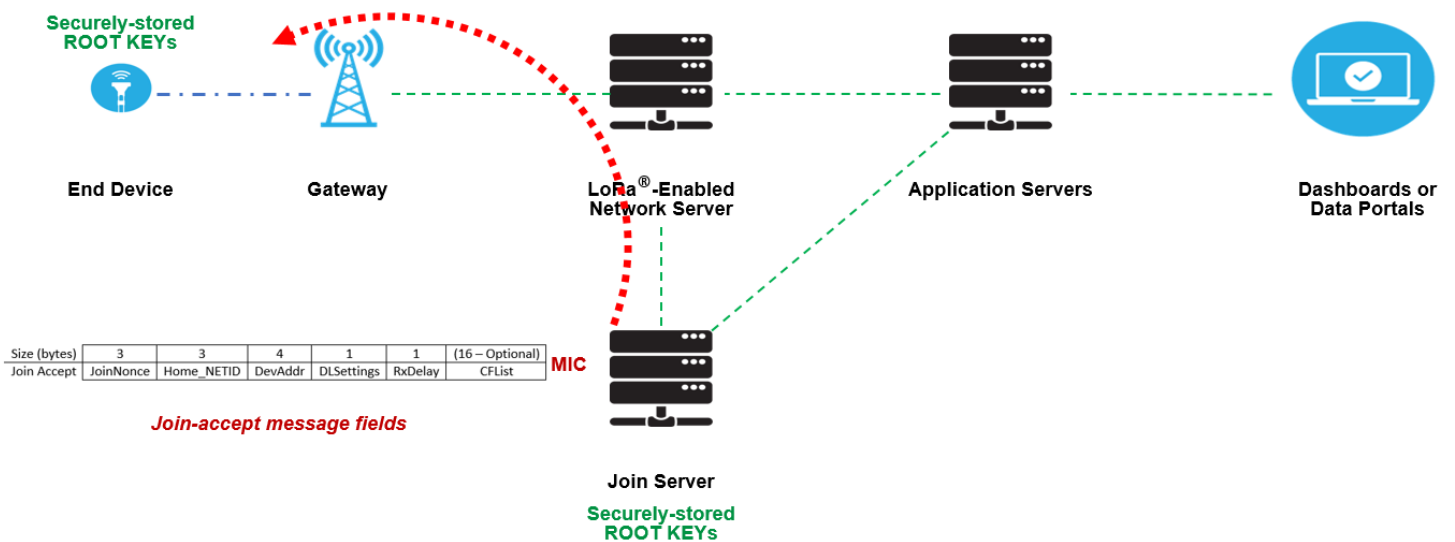


Figure 20. Sending a join accept message to an end device

Next, the end device **derives session keys locally**, based on the DevEUI, Join EUI, DevNonce, root keys and fields in the join request and join accept messages. On its end, the join server also derives session keys from the serial IDs, root keys and fields in join requests and join accept messages. Finally, the join server shares session keys with network and application servers, as illustrated in Figure 21.

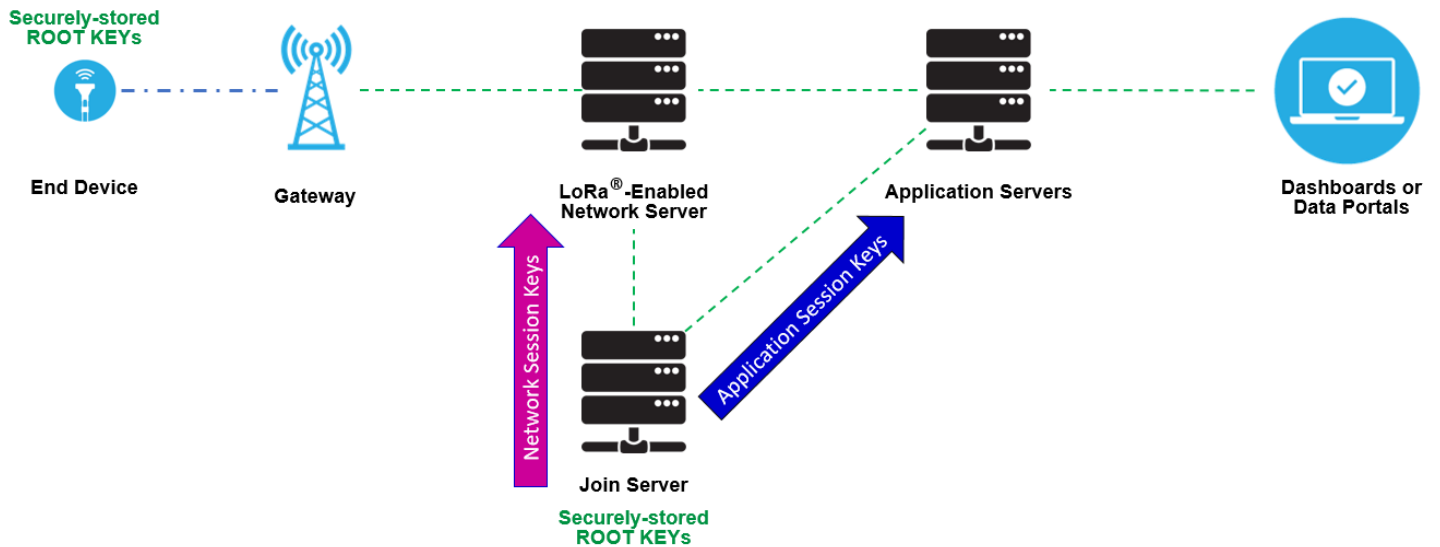


Figure 21. Session keys are shared with the network server and the application server

Figure 22 illustrates the security of data packet transmissions. The control traffic between the end device and the network server is secured with a 128-bit AES Network Session Key (NwkSKey). The data traffic that travels between the end device and the application server, is secured with a 128-bit Application Session Key (AppSKey). This method ensures that neither the gateway nor the network server can read the user data.

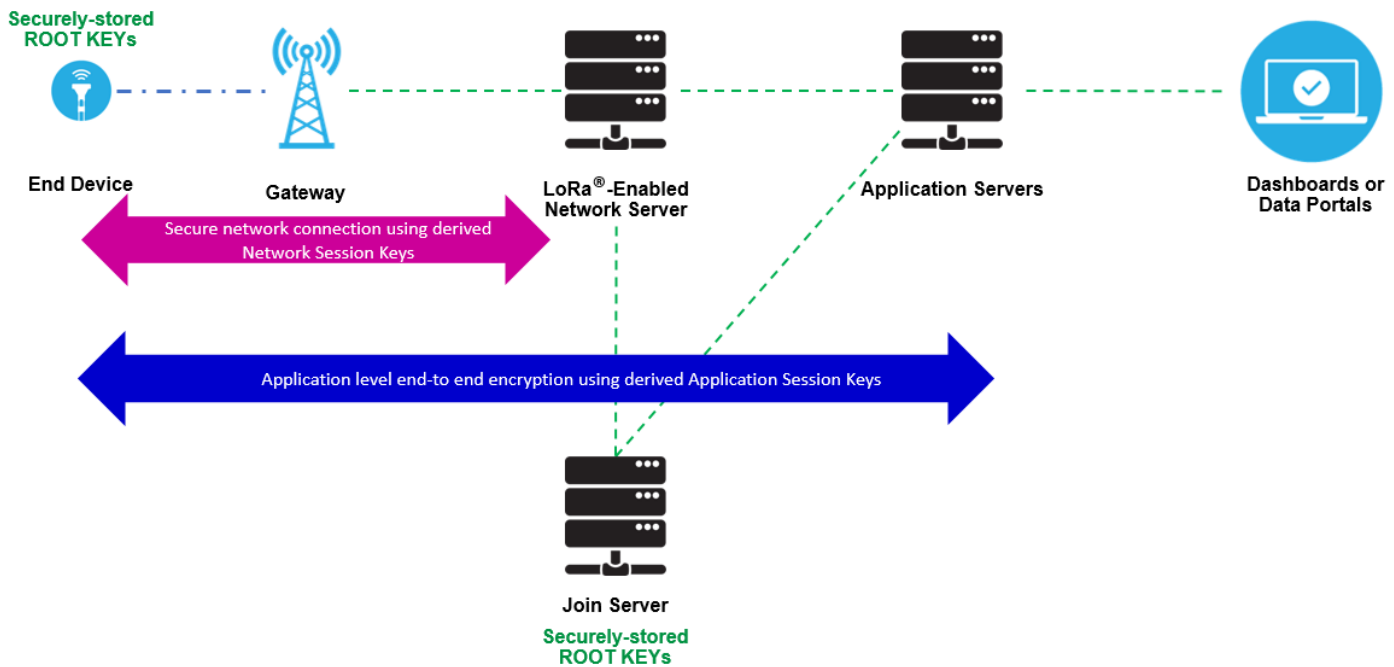


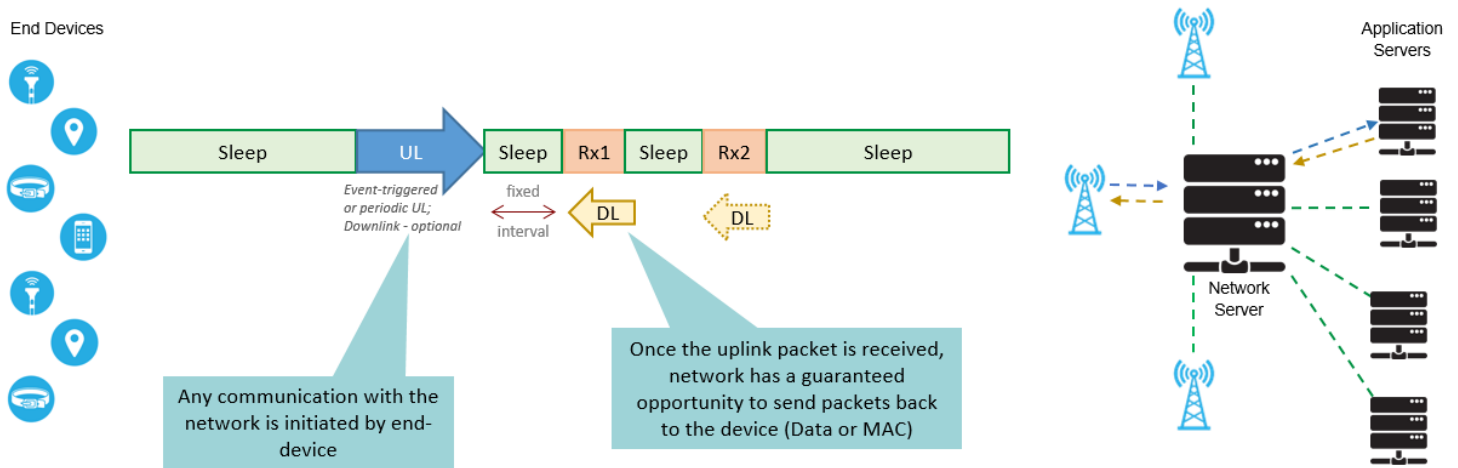
Figure 22. Secure transmission of data packets

3.4. Device Classes: A, B and C

LoRa-based end devices may operate in one of three modes, depending on their device class. All such devices must support Class A operation. Class B devices must support both Class A and Class B modes, and Class C devices must support Class A mode of operation with Class B mode being optional. These modes of operation have to do with how the devices communicate with the network.

3.4.1. Class A Devices

Figure 23 shows how the Class A mode of operation works.



There is no fixed association between an end device and a gateway. An uplink packet will be processed by all gateways within range and forwarded to the Network Server

Figure 23. Class A operation

In this case, the end device spends most of its time in an idle state, (that is, in sleep mode). When there is a change in the environment related to whatever the device is programmed to monitor, it wakes up and initiates an uplink, transmitting the data about the changed state back to the network (Tx). The device then listens for a response from the network, typically for one second (although this duration is configurable). If it does not receive a downlink during this *receive window* (Rx1), it briefly goes back to sleep, waking a moment later, again listening for a response (Rx2). If no response is received during this second Rx window, the device goes back to sleep until the next time it has data to report. The delay between Rx1 and Rx2 is configured in terms of a delay from the end of the uplink transmission.

NOTE

There is no way the application of the end device can wake up a Class A device. Given this limitation, Class A devices are generally not suitable for actuators.

Figure 24, Figure 25 and Figure 26 illustrate these communication patterns.

Receive Windows: Nothing is received

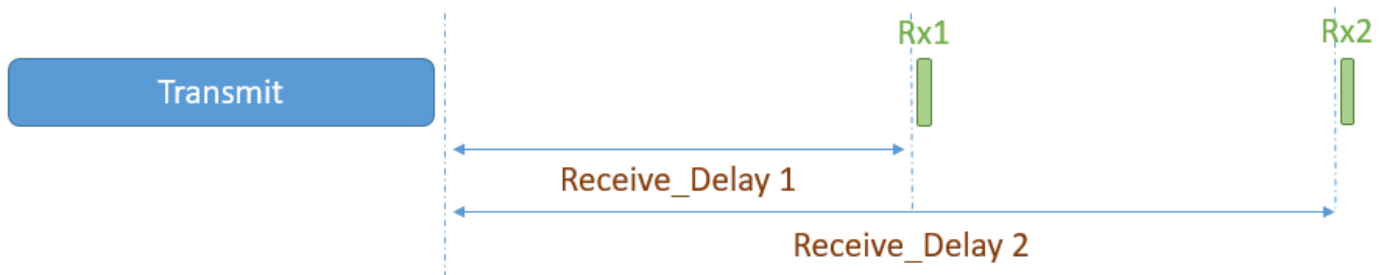


Figure 24. Class A operation when nothing is received

Receive Windows: Packet received in Rx1 window




Figure 25. Class A operation when a data packet is received in the first receive window

Receive Windows: Packet is received in Rx2 window




Figure 26. Class A operation when a data packet is received in the second receive window

NOTE

A device will not try to send another uplink message until either:

1. It has received a downlink message during Rx1, or
2. The second receive window following the last transmission is complete

3.4.2. Class B Devices

An enhancement of Class A, LoRaWAN Class B mode offers regularly-scheduled, fixed-time opportunities for an end device to receive downlinks from the network, making Class B end devices suitable for both monitoring sensors as well as actuators. All LoRa-based 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.

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.

Class B Beacons

For the Class B mode of communication to work, a process called *beaconing* is required. During the beaconing process, a time-synchronized beacon must be broadcast periodically by the network via the gateways, as illustrated in Figure 27. The end device must periodically receive one of these network beacons so that it can align its internal timing reference with the network.

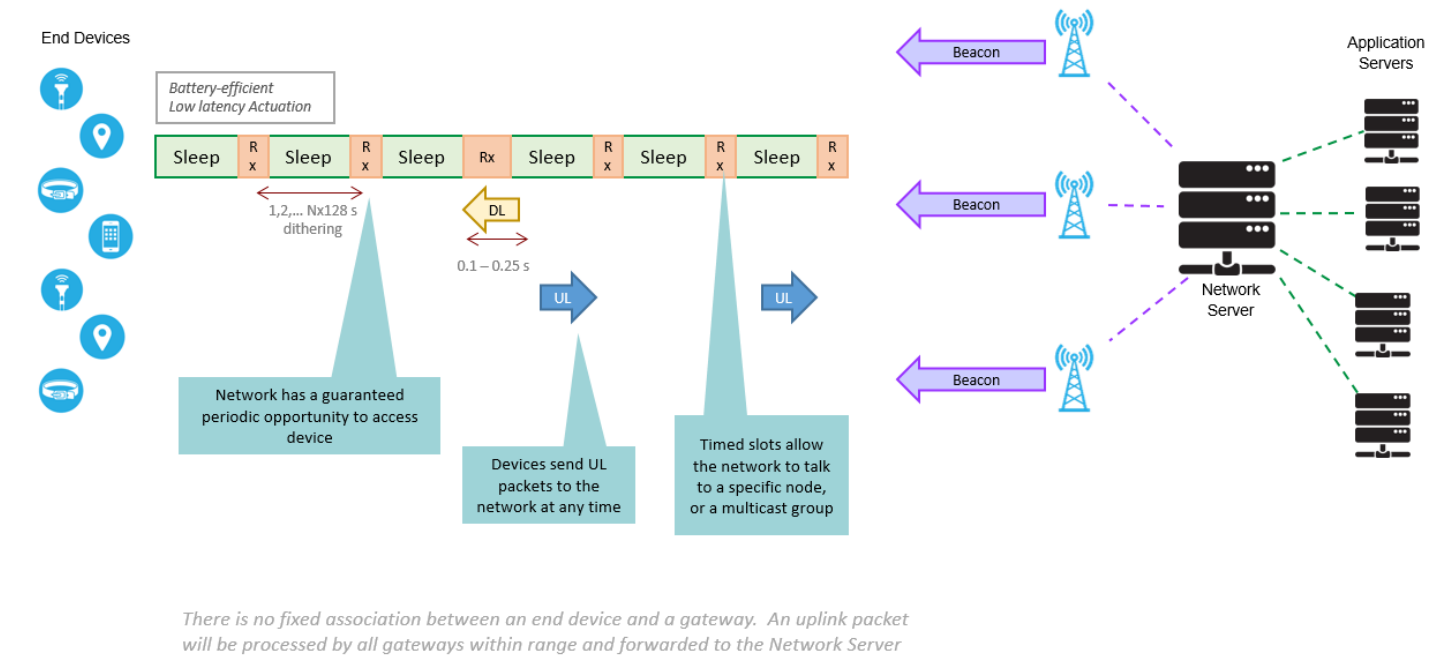


Figure 27. Class B beaconing operations

Devices use beacons to derive and align their internal clocks with the network. Devices do not need to process every beacon if the device is already aligned. In most cases, realigning several times a day is sufficient, with a minimal impact on battery life, as illustrated in Figure 28.

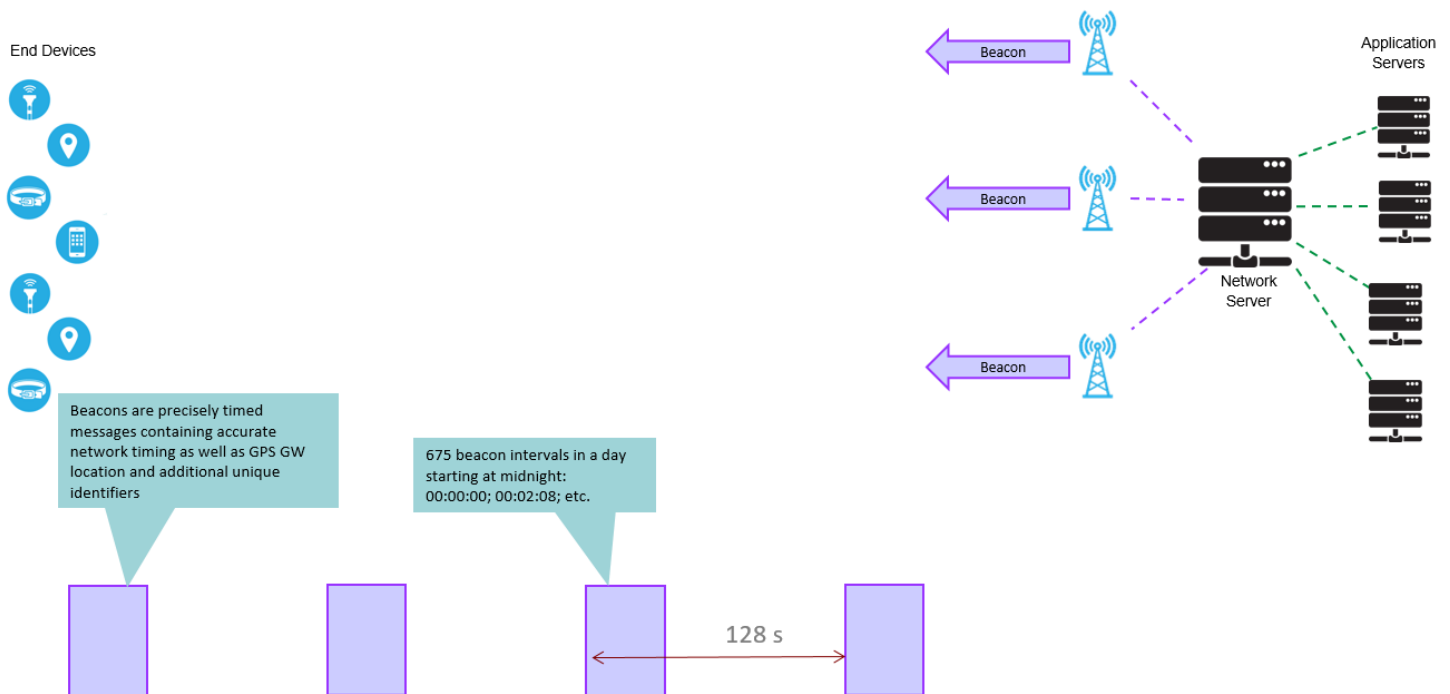


Figure 28. Periodic Class B beaconing for device synchronization

Based on the beacon's 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, as shown in Figure 29. In order for a LoRaWAN network to support Class B devices, all the LoRaWAN gateways in this network need to have a built-in GPS timing source, so they all can be synchronized to the exact beacon timing.




Figure 29. Class B ping slots

NOTE

Class B devices can also operate in Class A mode.

3.4.3. Class C Devices

Class C devices are always “on”; that is, they do not depend on battery power. Class C devices include such things as street lights, electrical meters etc. These devices are always listening for downlink messages, unless they are transmitting an uplink. As a result, they offer the lowest latency for communication from the server to an end device.

Class C end devices implement the same two receive windows as Class A devices, but they do not close the Rx2 window until they send the next transmission back to the server. Therefore, they can receive a downlink in the Rx2 window at almost any time. A short window at the Rx2 frequency and data rate is also opened between the end of the transmission and the beginning of the Rx1 receive window, as illustrated in [Figure 30](#).

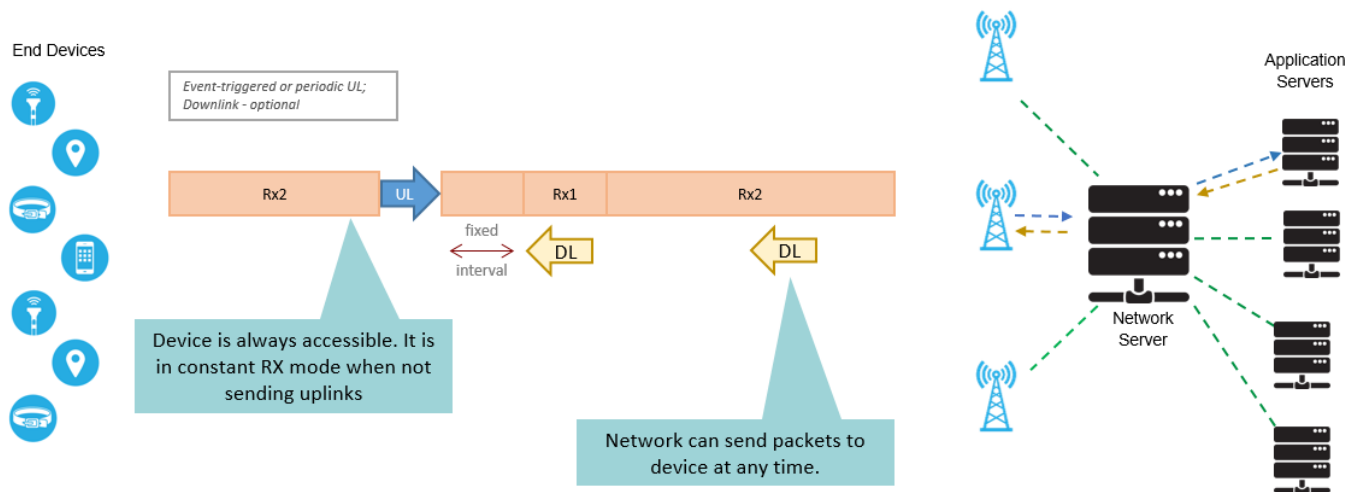


Figure 30. Class C operation

Chapter 4. The LoRa Alliance

With more than 500 member companies, the LoRa Alliance is one of the fastest-growing technology alliances. A community of innovators, the LoRa Alliance is committed to standardizing low power wide area networks (LPWANs). To this end, the group provides the LoRaWAN Specification (<https://lora-alliance.org/search/specification>) free of charge. The specification is based on open standards and provides for certified interoperability.

The LoRa Alliance also offers the [LoRaWAN Certification Test Tool](#), to help manufacturers ensure that their devices are fully LoRaWAN-compatible prior to sending those devices to an Authorized Test House for formal LoRaWAN Certification testing.

For more information on the LoRa Alliance, visit their website: <https://lora-alliance.org>.

Disclaimer

IMPORTANT

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 orders 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 customer 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 or guarantee, express or implied, regarding the suitability of its products for any particular purpose. All rights reserved.

© Semtech 2024