Capacity
QoS
Battery life
Range
Reliability
Standard
Security
Cost

WEIGHTLESS™

# LPWAN Technology Decisions: 17 critical features

## Introduction

Choosing connectivity technology for your next IoT project requires careful consideration of multiple technical and commercial factors, each intrinsically linked closely to your individual use case. In this paper we are going to discuss the criteria against which IoT connectivity technologies should be measured in order to help you make the right decision.

This is an initial report - we will be updating it periodically and will automatically send you revisions as they are published. You will also receive a series of detailed technical considerations of the key features of LPWAN technology options. This will put Weightless technology into context and provide a solid, objective basis on which to make a technology decision.

## Executive Summary

We first set out the IoT connectivity technology landscape from short range through LPWAN to 3GPP derived technologies operating in licensed spectrum. We then look in more detail at the several LPWAN options available comparing these objectively and making associated recommendations based on these. Finally we have reproduced a copy of the Table of Contents from the Weightless-P Open Standard offering an insight into the scope of the technology.

## IoT connectivity options

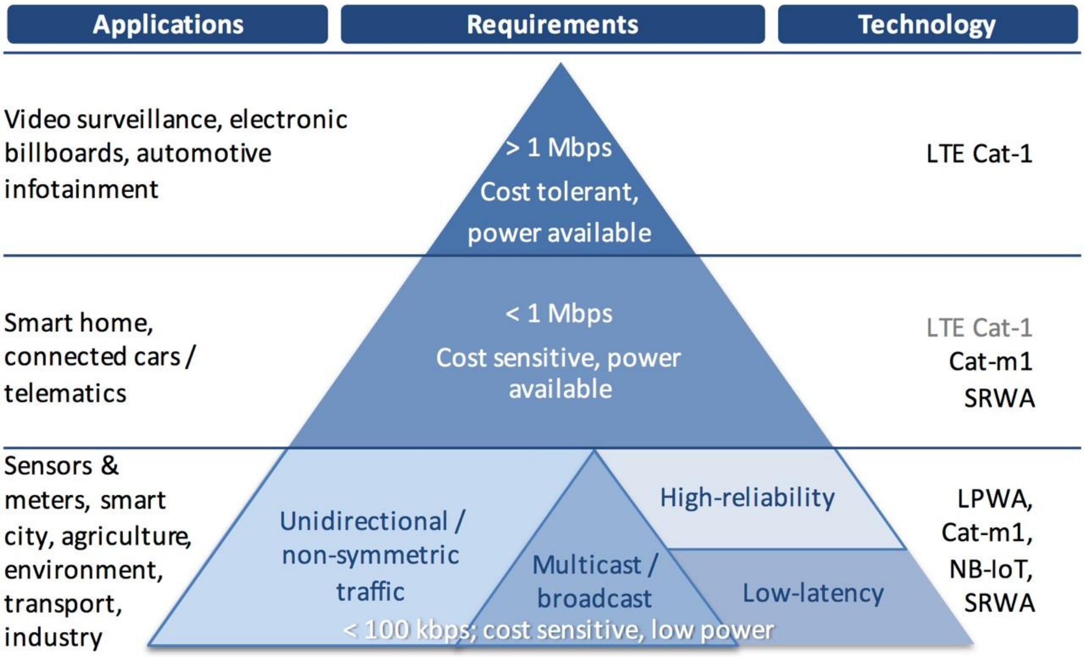
Applications
Requirements
Technology
Video surveillance, electronic billboards, automotive infotainment
> 1 Mbps
Cost tolerant, power available
LTE Cat-1
Smart home, connected cars / telematics
< 1 Mbps
Cost sensitive, power available
LTE Cat-1
Cat-m1
SRWA
Sensors & meters, smart city, agriculture, environment, transport, industry
Unidirectional / non-symmetric traffic
Multicast / broadcast
High-reliability
Low-latency
< 100 kbps; cost sensitive, low power
LPWA,
Cat-m1,
NB-IoT,
SRWA

Reproduced with kind permission. © 2016 Mobile Experts LLC. All Rights Reserved.

The range of options available to the IoT developer is considerable but every technology is subject to the same fundamental laws of physics with an inescapable relationship between range, data rate and transmit power. Greater range can be achieved by reducing data rate or increasing transmit power. A higher data rate means shorter range if transmit power is held steady. And a lower transmit power is possible if we can afford a shorter range or lower data rate. Ultimately the use case will determine which of these parameters a developer can afford to compromise and so help inform the decision on an appropriate technology.

Technologies such as Bluetooth, and in particularly Ble, Wi-Fi and Zigbee are ideally suited to use cases where range is not an issue and so low transmit power with relatively high data rates are possible. Typically these types of technologies will find applications in the home or office where a wider area network is not necessary to deliver the proposition.

At the other end of the spectrum where long range is required, perhaps with relatively high data rates, and where transmit power is less limited by regulations or a need for prolonged battery life LTE derived cellular technologies are common. NB-IOT technologies, due to reach the market over the course of the next few years, will typically service these high end use cases but at an inevitably higher cost and with higher power consumption due to more complex protocols, networks and higher transmit power.

Occupying the centre ground is a sector that has been emerging over the last few years and has quickly established itself as a key compromise between short range low power and long range cellular - low power wide area network or LPWAN technologies.

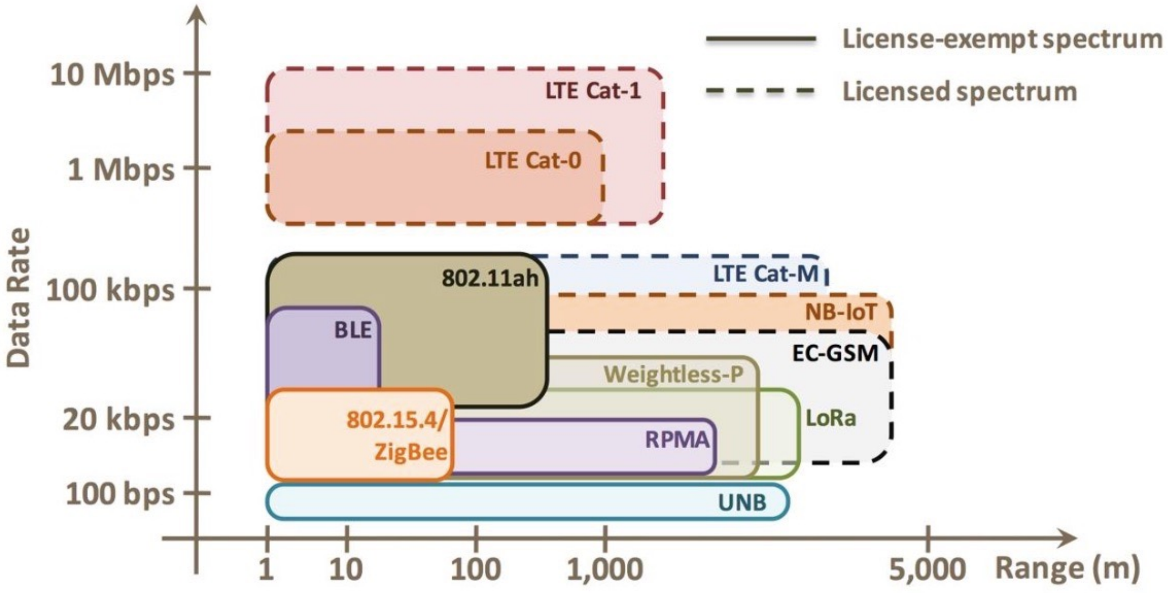

Reproduced with kind permission. © 2016 Mobile Experts LLC. All Rights Reserved.

## LPWAN

A number of companies have emerged to service the LPWAN space, each with propositions that fall on the spectrum roughly defined at one end as low cost and low performance through to near carrier grade at the upper end. The defining characteristic of this segment is that of modest data rates - consistent with the typical requirements of a very large proportion of IoT use cases. However, performance varies considerably across this spectrum.

Here we will compare and contrast three different technology choices as defined by the modulation schemes - ultra narrow band, narrowband and wide band.

## Ultra narrow band

The concept behind UNB is that as the transmission bandwidth is reduced the amount of noise entering the receiver also falls. With a very narrow bandwidth, the noise floor is reduced considerably, resulting in a large range for a low transmit power. However, an ultra-narrow channel can only carry low data rates and so UNB systems tend to be associated with small packet-size transmission. Also ultra narrow band cannot readily support bidirectional communications and so tends to be deployed in low end use cases where reliability and quality of service (QoS) are less important.

## Narrow band

As in many other situations, the optimum configuration or 'sweet spot' is often found between the two extremes and so in terms of QoS, capacity, cost and network efficiency, narrow band technologies offer a great compromise. Narrow band channels, around 12.5kHz wide, offer optimal capacity for uplink dominated traffic of moderate sized data payloads from a large number of terminal devices.

## Spread spectrum

An alternative approach is to adopt a wide channel – often 500kHz or over 1MHz – and then to use spreading of the data to gain range. This brings flexibility as the spreading factor can be varied depending on the channel conditions and in smaller cells very high data rates can be adopted. However, there are few spectrum bands wide enough to support multiple wide band channels and so different terminals and base stations need to share the same spectrum. Their transmissions can be differentiated through use of different spreading factors but only in a tightly time-synchronised and power controlled system with strong central control. This is difficult to achieve for the short and occasional communications made by most IoT terminals and rules out multiple public and private networks.

## What matters?

Choosing connectivity is complex - there are a lot of features and benefits, often conflicting, to weigh in the balance. So let's distil some of the key characteristics that will define an IoT connectivity technology and from there we can more easily make decisions about what is important to our particular use cases. We think that the strength of an IoT connectivity technology can be defined in terms of the following eight parameters.

- Capacity
- Quality of Service
- Range
- Reliability
- Battery life
- Security
- Cost
- Proprietary vs Standard

Below we detail how each of these parameters can impact on your IoT project. We then list a number of characteristics that show how Weightless technology addresses the requirements in the context of these parameters.

Please note that we are preparing detailed technical information about the core competitive advantages of Weightless-P. We will be making this available to subscribers to this paper in due course. We are also preparing a detailed, interactive, technical webinar programme. This is intended to answer any questions that you have about Weightless-P technology to help you make an informed choice about IoT connectivity design.

# CAPACITY

## Excellent capacity and scalability for IoT deployment

Today, the relatively few LPWAN networks in existence are generally not capacity constrained. There are currently few devices connected to them, often in trial mode, where aspects such as range and functionality are being tested. But if the predictions of 50 billion devices are even remotely correct then this will all change in the next few years as the number of devices grows rapidly. Base stations could have hundreds of thousands of devices connected at any point and just as the iPhone stimulated the data crunch, we could see a "machine crunch" on the emerging networks.

We are used to discussing and defining capacity on cellular networks. This is typically stated in measures such as bits/Hz where a technology is rated according to how much data it can transfer per unit of radio spectrum owned by the operator. Each generation strives for better efficiency through improved radio technology, better scheduling of traffic and so on. When technology is unable to cope operators deploy a mix of additional spectrum where they can acquire it through auction and smaller cells. The latter has been the key driver of capacity growth over the recent decades, delivering much of the 100-fold or so increase in capacity needed to meet smartphone requirements.

When designing for IoT it is second-nature for engineers to reach into the same toolbox as for cellular. But this is a mistake – there is much about IoT that is very different and many of the same techniques just do not work. Some of the key differences are:

- Short messages. Many IoT devices send tiny amounts of data. A car park sensor need only send 1 bit – whether the space is full or empty. A thermostat might need 8-16 bits. A locating device perhaps as much as 8 bytes. These volumes are often swamped by the packet size. For example, if IPV6 addressing is used then an address is 128bits long. A device reporting its identity then its message could increase data volumes 10-fold. Other signalling messages such as location updates could similarly consume a large amount of network resources – a device in a vehicle that updates its location for network management purposes but only transmits once a day could send 1,000-times more data than the end user wants. Designing carefully for short messages could easily improve capacity 10-fold.

- Random timing. Most cellphone interactions start at a random time – the point when someone wants to call a user, or they decide to instigate a search. The device then goes through a "random access" phase to initiate communications with the network after which the network provides dedicated resource for the duration of its communications or "session". Random access is relatively inefficient. There is a chance that multiple users try to access the network resource at the same time and clash. When this happens often all communications are lost and the users have to repeat their transmission hoping that it does not clash second time around. The efficiency of such channels is well defined in the "Aloha access" theory which tells us that, at best, they can be about 30% used. Above these levels there are so many message collisions and re-transmissions which then re-collide that the channel capacity spirals downwards and a reset is needed. In the cellular world the random access phase is only a tiny, tiny fraction of the total data transmitted and so its inefficiency is of little relevance. In the IoT all of the data can be encapsulated in the first message meaning that all transmissions could be random access. In this case efficiency drops to 1/3 at best. If devices could be told when to transmit next – for example thermostats given periodic slots – then 3-fold efficiency improvements can be made.

- Power adjustment. In cellular systems handsets are tightly controlled by the network to use the optimal type of modulation and power levels, with these varying often second-by-second. In an IoT network transmissions are so short that there is insufficient time for the network to adjust the device. Hence, the device tends to use higher powers than needed resulting in more interference. Networks need to be designed both with intelligent ways to adjust device power based on knowledge such as whether the device is static (and so transmit power can be steadily adjusted over time) and other cues from the network.

- Multiple overlapping networks. Cellular operators have their own spectrum and can design networks free from interference from others. Conversely, most IoT networks are deployed in unlicensed spectrum where there can be interference from other IoT networks using the same technology, other IoT networks using different technology and other users. To date, this has not been a key issue but as competition grows and more networks are deployed it is very likely to become a constraining factor. Some techniques, such as code-division access (CDMA and similar) rely on orthogonality between users which is only effective where users are controlled in time and power. With a single network this is possible, but with multiple networks there is rarely coordination between them and the impact of interference can be severe. Instead, techniques such as frequency hopping and message acknowledgements are much more important as are networks that can adapt to their interference environment.

- Broadcast messages. Many applications require the same information to be sent to all devices. This might be a software update, new tariff information, share prices and so on. If there are 10,000 devices in a cell then sending this information separately to each is extremely inefficient. Broadcast capabilities (to all devices in a cell) and multicast (to a selected sub-set of devices) are critical in ensuring efficient operation with the ability to provide 10-fold efficiency gains or higher depending on message frequency and size.

- Alert messages. Often devices will be triggered by a common event into sending an alert. For example, a power outage in a city might result in all smart meters simultaneously trying to send a "no power" message. This immediately congests the network often resulting in network failure. Means to detect the first message and then tell all the meters to discard their alarm message can help the network avoid these overload situations. This then means the network can be sized for lower peaks in demand and so run more efficiently.

For all of these reasons and more, the efficiency of an IoT network should not be measured in the classical manner. A network could have apparently worse modulation but simply through smaller message sizing be 10-times more efficient.

There are many technologies that are sub-optimal and have the potential to suffer severe capacity constraints. For example, the Sigfox solution retransmits messages multiple times which is clearly inefficient and has limited ability to take any action once a cell is overloaded. The LoRa system relies on orthogonality between transmissions which could suffer badly when multiple overlapping networks are deployed in the same spectrum. Cellular solutions are still in definition but often have large minimum packet sizes. These issues may not become apparent during a trial where network capacity is not stressed and only emerge when tens of thousands of devices are deployed. At this point, changing the technology is very expensive.

Weightless has been carefully designed from the start for the 50-billion device world. It has very short message sizes, frequency hopping, adaptable radios, group messages and multicast messages, minimal use of random access through flexible scheduling and much more. Although its bits/Hz may not be materially different from other solutions, in practical situations it is potentially orders of magnitude more efficient. If IoT devices were like handsets are replaced every two years, that might not matter, but with some being 10-20 year deployments, getting it right now is critical.

When IoT connectivity technologies are being considered, adopters rightly consider parameters like cost, battery life and range but it's easy to overlook the importance of network capacity. And in the absence of empirical data from real world network deployments, we can be tempted to rely too heavily on theory rather than properly modelled scenarios ignoring the critical parameters limiting network capacity. Capacity is not just about the number of simultaneously connected nodes, it is about mean data packet length, transmission time, frequency of transmissions and interference mitigation. Scaleability is also about the fundamental technology behind the network from ultra narrow band at one end of the spectrum to spread spectrum at the other. Not surprisingly, both give rise to upsides and downsides and the capacity sweet spot is a compromise between the two. Narrowband channels offer the optimal capacity for uplink dominated traffic from a large number of devices with moderate payload sizes and typical transmission duty cycles so Weightless-P uses 12.5kHz channels. Weightless-P also offers flexible channel assignment which further enhances network capacity by enabling frequency reuse in large scale deployments and adaptive data rates from 200bps to 100kbps permit optimal radio resource usage to maximise capacity. Time synchronised base stations allow for radio resource scheduling and utilisation. Weightless-P operates across the entire range of licence exempt sub-GHz ISM/SRD bands for global deployment. Let's now consider a real network scenario and compare the alternative technology options.

Normally we think of network capacity as a measure of the number of end devices connected simultaneously to a single base station so we'll keep with this convention and see how MAC total throughput, transmission frequency and data payload all factor into a capacity calculation below.

Figure 1 shows the MAC throughput for three popular LPWAN technologies based on ultra narrow band, spread spectrum and, for Weightless-P, narrow band. These numbers pertain to EU regulations; the data rate under FCC regulations is different because of higher Tx power but narrow band continues to offer an order of magnitude higher MAC throughput in this region.

Figure 1

| technology: | ultra narrow band | spread spectrum | narrow band |
|---|---|---|---|
| MAC throughput bits/s | 1404 | 93 | 4923 |

How did we calculate this?

We have made objective assumptions on the modelling criteria. In fact we have been considerably more conservative in our assumptions than other published models suggest for well known UNB and spread spectrum technologies. And we declare the criteria - here it is:

- Weightless-P adaptive data rate with 10dB margin, PER target 0.1%
- Scheduled up-link capacity
- Mean data rate determined by throughput for randomly positioned nodes with properly assigned data rate
- Weightless-P: -134dBm sensitivity for 0.625kbps, EU Tx power is 14dBm and US Tx power is 27dBm
- UNB and narrow band MAC throughput based on urban Hata model (BST antenna height 30m and ED height 0.5m)
- Coverage for narrow band is 1.5km for EU and 3.8km for US
- spread spectrum MAC throughput based on Ingenu white paper but adjusted for over-conservative assumptions on repetition rate
- Capacity loss due to slot granularity is absorbed by assuming 50% protocol overhead and 50% UL half duplex ratio

In any wireless system the data throughput determines the achievable network capacity. Higher data throughput enables larger data packets, more frequent transmissions and a greater number of end points. These fundamental parameters are the key factors in the scalability debate. Increase any one of these and you are stress testing the scalability of the network. Let's look at a typical scenario.

Smart Meters

In the utility metering sector a 15 minute reading interval is the accepted default frequency of uplink transmissions. And a data packet of 200 bytes would be considered normal. What does this mean for an ultra narrow band, spread spectrum and narrow band technology?

200 bytes every 15 minutes is 800 bytes/hour or, alternatively, 1.78 bits per second. The MAC throughput divided by the end device data throughput will define for us the number of nodes that can be serviced - this is how data rate and capacity are linked.

Weightless-P can handle 2769 end points per base station with these uplink characteristics. A spread spectrum technology can manage 52 and an ultra narrow band technology can accommodate 789 end points.

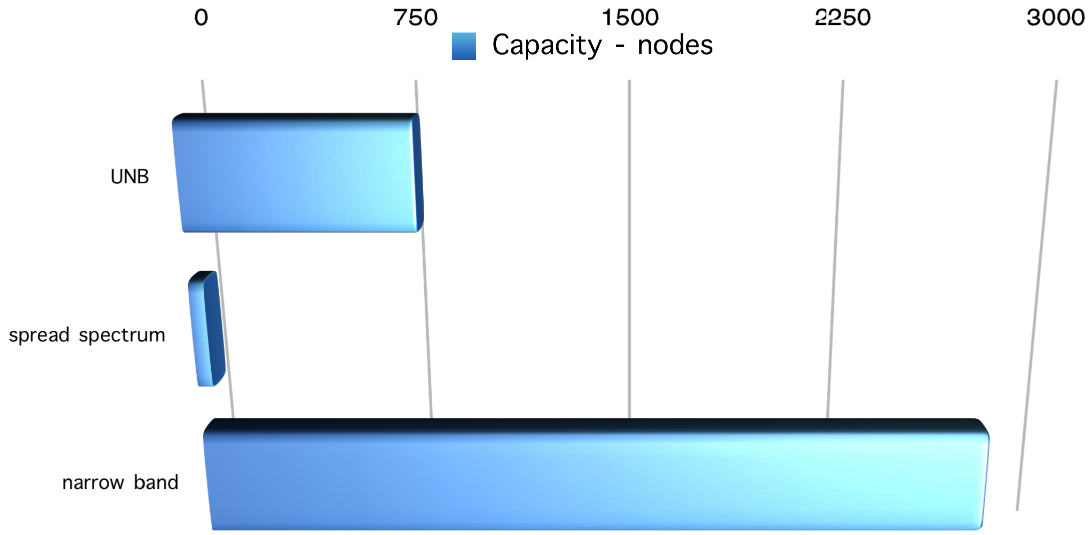

.

## What Weightless offers:

- FDMA+TDMA in 12.5kHz narrow band channels offer optimal capacity for uplink-dominated traffic from a very large number of devices with moderate payload sizes
- Operates over the whole range of license-exempt sub-GHz ISM/SRD bands for global deployment: 169/433/470/780/868/915/923MHz
- Flexible channel assignment for frequency re-use in large-scale deployments
- Adaptive data rate from 200bps to 100kbps to optimise radio resource usage depending on device link quality
- Transmit power control for both downlink and uplink to reduce interference and maximise network capacity
- Time-synchronised base stations for efficient radio resource scheduling and utilisation

# QUALITY OF SERVICE

## Fully acknowledged transmissions

When IoT users talk about quality of service (QoS) they are typically referring to particular requirements for message delivery such as speed of arrival and probability of successful reception. This might be encapsulated in statements such as "emergency messages much be received within 2s".

Most users understand that there are trade-offs associated with QoS. For example, message delay is often caused by network congestion, so simply over-sizing the network can improve QoS, but at a cost. Short delay times typically limit the ability of devices to move into idle mode and therefore result in much shorter battery life. A network designed just for high QoS requirements would be a little like a parallel road system designed just for emergency vehicles – it would do the job well but be prohibitively expensive.

The best approach is a flexible system that can deliver different QoS outcomes according to need, prioritising the high QoS traffic over the low. This is more akin to our current road system where vehicles with flashing blue lights can move ahead of others in any congested situation.

Before looking at how best this might be achieved, it is worth exploring the limits of QoS guarantees. No system can provide an absolute guarantee that all high QoS messages will be delivered successfully in a short time period. The device may be out of coverage. A base station might be down. There might be local interference from other users. The best that any network can do is aim for a high probability (eg 99%) that under specified use cases the requirement will be met.

It is sometimes claimed that only networks using licensed spectrum can reliably provide high QoS. This is predicated on the assumption that in unlicensed spectrum interference cannot be controlled and without full control of all relevant factors the operator cannot provide the guarantees needed. Clearly having more control over all key factors makes it easier to meet requirements, but it is not necessary. To continue the analogy, parcel delivery companies guarantee next-day QoS even though they use shared road systems rather than their own dedicated network. They succeed because the level of congestion is generally predictable and systems in place to mitigate the effects of unexpected congestion such as re-routing solutions. The same is true in unlicensed spectrum. Interference will tend to build slowly over time and be predicable and increasingly "mapped". Solutions such as frequency hopping can mitigate the worse effects and in the longer term interfering users could opt to coordinate between themselves or additional frequency bands could be added to the solution. It may require the operator to work a little harder, but experience suggests that dedicated spectrum is not a prerequisite for providing high QoS.

High QoS messages are nearly always unpredictable as to when they will occur (the routine report of "all's well" from a monitoring unit rarely requires high QoS). If the messages originate at the device then there needs to be a rapid and reliable random access mechanism. One of the best ways to achieve this is to have a subset of the random access resource set aside for high QoS applications. This subset can be heavily over-dimensioned without having a major impact on the rest of the network. However, this does require that the system specifications detail how the sub-channel will work, where it will be located and under what circumstances devices are entitled to use it. If the messages originate in the network then there is less difficulty in getting the network to schedule them immediately but devices need to be awake to receive them. This means that these devices need to be very frequently monitoring paging channels which will inevitably result in increased battery drain. By allowing devices to determine their optimal idle time, such "always-on" devices can be accommodated without impacting the rest of the network.

There are other factors that might also come into play. The times when high QoS is required are often those when a major incident occurs such as a terrorist attack. Such an incident can trigger communications from all, congesting networks. This may be equally true in IoT networks where a power outage can trigger all smart meters in a city to send repeated "power out" alerts. This can swamp networks to the extent that a re-boot is needed resulting in no messages getting through until the network has overcome the congestion. Having solutions that can deal effectively and quickly with widespread alert messages can ensure networks stay available.

Ultimately, affordable QoS is all about intelligent network design. This needs to be at the heart of the technology standard with a range of options to meet diverse needs.

LPWAN technologies fall into three categories - uplink only, uplink and partial downlink capability and 100% fully acknowledged transmissions. Weightless-P supports full acknowledgement of all transmissions where required. Some alternative LPWAN technologies offer downlink capabilities limited to a very small proportion of transmissions meaning that reliability and QoS is compromised. Weightless-P also supports acknowledged and unacknowledged unicast and multicast traffic. It offers a flexible acknowledgement scheme including deferred and combined acknowledgement for improved resource usage. It also supports both network originated and device originated traffic with paging capacity and low latency in both uplink and downlink. It enables fast network acquisition, Forward Error Correction (FEC), Automatic Retransmission Request (ARQ), Adaptive Channel Coding (ACC), handover, roaming and cell re-selection. Real bidirectional capability also supports over-the-air firmware upgrade and security key negotiation or replacement.

## What Weightless offers:

- Supports both network-originated and device-originated traffic
- Paging capability
- Low latency in both uplink and downlink
- Fast network acquisition
- Forward Error Correction (FEC)

.

- Automatic Retransmission Request (ARQ)
- Adaptive Channel Coding (ACC)
- Handover
- Roaming
- Cell re-selection

# RANGE

## Long range in urban environment

Range is one of those factors that all technologies are keen to discuss, but not something directly needed by the end user. Users of IoT devices want connectivity at the lowest possible cost. Long range can assist with this by reducing the number of base stations an operator needs to deploy in order to provide the required coverage.

However, range is not the only factor impacting the number of base stations as more may be needed to provide sufficient capacity. This is likely to become increasingly true as the number of IoT devices increases. In cellular systems range is rarely discussed these days because most cells are now smaller than the maximum in order to deliver the needed capacity. So a solution that delivered long range but low capacity might end up requiring more base stations than a more balanced system.

Range may also be of less relevance in some cases such as smart cities. Here, city wide coverage can often be achieved with a handful of base stations, and reducing this by one or two because of greater range does not make any material difference. It is only where rural coverage is required that longer range becomes important.

Quoting range is a little like quoting car fuel consumption. Unless the test conditions are identical, comparing two cars, or technologies, will not be meaningful. Range is impacted by many factors from antenna height, to terrain, to frequency, to interference levels. A better approach is to quote the maximum path loss that can be tolerated which is more comparable. In practice, most competing systems will have similar parameters. All systems have the same constraints in terms of transmit power allowed by the regulator and noise floor imposed by physics. All can make the same trade-offs of lower data rate for greater range using techniques such as spreading or ultra-narrowband emissions. Most will allow for flexibility so different terminals can select different combinations of range and data rate. The only variable that makes a material difference is the bandwidth with systems deployed in broad bands having more bandwidth they can trade against range. This is why TV white space appeared so attractive, but unfortunately failed to become widely available.

So it is reasonable to assume that in practice all well-designed LPWAN technologies will have a similar range. The number of base stations is then more constrained by capacity than anything else. This is the factor that those evaluating a technology should concentrate more on.

Range is frequently cited as one of the key parameters of an IoT system, especially in challenging conditions such as dense urban environments and in locations where the antenna position is compromised such as inside a building. LPWAN technologies are commonly promoted with varying claims with respect to range but the reality is that all technologies operating in sub-GHz unlicensed spectrum are subject to the same conditions, regulations and laws of physics. Range is ultimately determined by signal path, link budget, antenna size, quality, position and location, data rate and transmission power. Lower data rates with channel coding provide for a similar link budget to alternative LPWAN technologies and so achieve a typical range of 2km in an urban environment. In reality it is only possible to offer a typical range - a more accurate claim would require actual modelling in the specific environment. Weightless-P is specified with a realistic range in a dense urban environment which is closer to that achievable from all licence exempt LPWAN technologies.

## What Weightless offers:

- Lower data rates with channel coding provide similar link budget to other LPWAN technologies
- 2km in urban environment

# RELIABILITY

## Industrial-grade reliability

By reliability we normally mean messages getting through consistently day after day, throughout the lifetime of the device. There is also an element of QoS in that the messages should be received within a desired time frame – this is covered in a separate note [link].

There are a number of factors that could prevent message transmissions including:

- Hardware failure.
- Software bugs that either become worse over time or cause issues such as re-starts.
- Growing network congestion.
- Worsening reception perhaps due to increased interference or deteriorating antenna.

Most important in delivering a reliable network is acknowledgement of messages. A device, or the network will then know if a message fails to get through, or takes longer than anticipated and can flag this issue to users or network managers. This enables problems to be resolved quickly and deteriorations in performance to be spotted and diagnosed. It also allows devices to re-transmit messages until they do get through, helping to provide a reliable service even in unfavourable conditions.

Software bugs, glitches or security loopholes can occur at any time. We've become accustomed to weekly updates to Microsoft Windows and Android Apps. While IoT devices should need far less frequent updates, they will nevertheless need to have software changes from time to time as networks and services evolve. Only those systems that can efficiently and quickly upgrade the device software will maintain reliability across the years.

Network congestion can be overcome in many ways, but one of the best approaches is to have a highly efficient system such that congestion may never occur, and when it does it results in graceful degradation rather than system failure.

Choosing a LPWAN technology has traditionally involved a compromise - high performance, high cost and high power consumption from cellular based technologies operating in licensed spectrum or more modest performance from a low cost, low power technology in licence exempt bands. Today the decision will typically be driven by the use case; where cost and power consumption are not critical and where Quality of Service and reliability are priorities then the larger amounts of licensed spectrum with less restriction on transmit power suggests 3GPP technologies. But Weightless-P represents a game changer with its ‘clean slate’ design philosophy. By leveraging proven concepts from cellular technologies with a robust and tailored MAC and PHY, Weightless-P brings carrier grade performance at a low cost, low complexity price point. Indeed, Weightless-P supports licensed spectrum operation meeting mandatory selectivity/blocking performance requirements ensuring good coexistence in licensed bands without any protocol changes. Across all spectrums, including sub-GHz bands, Weightless-P offers industrial grade reliability.

## What Weightless offers:

- Fully acknowledged communications
- Auto-retransmission upon failure
- Frequency and time synchronisation
- Supports narrowband channels (12.5KHz) with frequency hopping for robustness to multi-path and narrowband interference
- Channel coding
- Supports licensed spectrum operation

# POWER

## Ultra-low energy consumption

Many IoT devices are battery powered, with batteries that are intended to last for many years. For such a device each bit transmitted is a bit closer to death! The battery life is impacted by both the power consumed during downtime (or "idle") and the power consumed when transmitting data.

The length of time a device spends idle is often a compromise between longer battery life and the ability to contact or page a device when required. The optimal balance will depend on the application – for some a daily chance to communicate with a device will be sufficient, for others there may be a need to contact a device in just a few seconds. The best systems do not impose any particular idle time but allow the devices to make the decision themselves based on pre-programmed or downloaded information.

Once out of idle mode and checking paging channels it is important to minimise receive time. Some solutions require devices to listen to an entire system information transmission and paging channel to check whether anything has changed. Others are more intelligent, providing flags at the start of the reception process which help the device determine whether it is worth listening to the remainder of the frame. Having multiple paging sub-channels also means the device only has to listen to one, reducing reception time.

Transmitting is hugely more energy intensive than receiving. Large gains in power consumption can be made with careful design. There are obvious steps such as selecting modulation formats that do not require linear power amplifiers and hence enable operation in efficient parts of the power curve. But far more important is to reduce the number of bits that need to be transmitted. Essentially, all systems face the same amount of path loss and overcome this with various mechanisms that lengthen the transmit time per bit such as the low data rates in ultra-narrowband or the multiple spreading code bits sent per bit in spread spectrum solutions. So broadly they all use about the same amount of energy per bit. Where they differ markedly is the number of bits. For example, some solutions automatically transmit the same message a certain number of times (eg three) to increase the likelihood of successful reception. Obviously this triples the power drain for transmission. Others require longer packets with large addresses or other overheads. This can easily incur a 10-fold increase in the number of bits transmitted compared to a highly optimised solution.

It is relatively easy for any LPWAN solution to claim 10+ year battery life – and if a device has prolonged idle periods and hardly transmits anything this can likely be achieved. But we all have the personal experience of how manufacturers' claims for battery life can often be hugely overstated in the real world and how very low power devices can be annoying slow to respond. It is difficult at this stage to verify whether a 10-year battery life under "typical conditions" can actually be met – until networks are loaded to near capacity and devices have been in the field for some time there will be limited evidence to go by. But finding out halfway through the anticipated device lifetime that 50 million units all need an unexpected battery replacement will be costly and disruptive.

When picking an LPWAN technology is it worth asking whether the designer really understood how to deliver real-world long battery life. Did they minimise the number of bits transmitted? Have they allowed the device to optimise its balance between battery life and the ability to be contacted? Have they selected radio parameters predominantly on the basis of their ability to be energy efficient using non-linear power amplifiers? Have they avoided blindly repeating messages or long address fields? Weightless designers certainly took all these factors into careful account.

Many IoT applications terminals will be powered by batteries leading to the need for low energy consumption. All LPWAN technologies operating in licence exempt spectrum are low power but Weightless-P uses a number of techniques to offer best in class power consumption performance and consequently very long battery life. Overall energy consumption is a factor of both the power used during transmission and the amount of the time that the transmitter is active - the duty cycle. Weightless-P uses GMSK and offset-QPSK modulation schemes which deliver optimum power amplifier efficiency. Offset-QPSK modulation is also inherently interference immune and using Spread Spectrum for improved link quality in busy radio environments minimises the required transmit power. A limit of 17dBm in ISM spectrum means that terminals can operate from coin cell batteries. Adaptive data rate also permits minimal transmit power for nodes with a cleaner signal path to the base station so maximising battery life. And since a terminal device will nearly always spend a very large proportion of the time in an idle state, the power consumption in this mode becomes critical. Weightless-P consumes less than 100$\mu$W when inactive.

## What Weightless offers:

- GMSK and offset-QPSK modulation for optimal power amplifier efficiency
- Interference-immune offset-QPSK modulation using Spread Spectrum for improved link quality in busy radio environments
- Transmit power up to 17dBm to allow operation from coin cell batteries
- Adaptive transmit power and data rate to maximise battery-life
- Power consumption in idle state when stationary below 100$\mu$W

# SECURITY

## Unbeatable secure networking

Mention IoT and very quickly the issues of security and privacy will be raised. The two are somewhat related. If data is kept secure then the companies that can read it are limited and privacy more likely to occur. If the system is not secure then the chances of achieving privacy are low.

We are so used to hearing about security breaches in Internet-related systems that it is easy to assume that nothing is secure and that any system can be hacked. This is not entirely true. Wireless systems can be very secure – there have been few, if any, significant breaches of the cellular systems that have resulted in customer conversations being overheard or account information hacked. When breaches do occur they tend to be at the application layer with organisations like Facebook and others struggling to prevent loss of password data or similar. So a well-designed wireless solution should be able to keep data and privacy secure up to the point that the data passes to a solutions provider (such as a car manufacturer). After this point, it is out of the control of the wireless solution as to what happens to the data. Equally, a poorly designed wireless system can provide opportunities for attack that are hard to close off since updating remote terminals can be difficult, especially in one-way networks.

There are many different elements of security including:

- The network authenticating the terminal to be sure that the terminal is the device it claims to be.
- The terminal authenticating the network to be sure it is a valid network to which it can pass information.
- Encryption of the information such that it cannot be overheard.
- Prevention of replay-type attacks where data is recorded and then replayed later to the network resulting in what appears to be the same message being resent.

The level of security selected is a compromise between many factors including:

- Ease of commissioning the device for the first time – ideally not requiring the commissioning engineer to enter long digit strings into databases or similar.
- The overhead added to messages to provide the encryption which should not be so large it materially increases data traffic volumes.
- The processing power required in the terminal to perform any security-related operations which ideally should not require additional elements or higher power drain.

- The power of encryption which ideally should not prevent export worldwide.

A key choice is whether to embed a secret key within the terminal. This is the approach adopted by cellular systems where each SIM card has a secret key inaccessible to anyone once fabricated and brings many advantages since the terminal sets out from the factory with all the information within it to enable the network to authenticate it. However, it does require a secure database to be administered among chipset manufacturers and network operators and tends to work best in an open standards environment. It is the approach adopted by Weightless which enables it to achieve "carrier grade security".

Another important design feature is the ability to upgrade the entire security suite over the air. This means that should a flaw be discovered a new security approach can be downloaded and installed remotely, resolving the problem. Again Weightless offers this capability.

There are other subtleties with IoT. Devices often send the same message repeatedly such as a meter reading or similar. Sending the same message generates a security weakness that attackers can use to decode (it was how the Enigma cracked the German code – the use of the phrase "Heil Hitler" at the start of many messages). Weightless overcomes this by generating a changing number called a "nonce" that is encoded along with the data from the device to ensure the message is always different. This also allows various other security checks such as that messages are arriving in sequence, preventing messages being recorded and replayed later.

So an IoT system can be made extremely secure – as secure as a cellular solution. At this point the weaknesses are much more likely to be with the client-stored data than the wireless network. But it requires excellent design, careful selection of trade-offs and a belt-and-braces approach of being able to swap out a security suite that turns out to be weaker than expected. Unfortunately, few systems available to date have all of these.

It only takes one widely publicised breach to remove all confidence in a technology. Weightless is designed to ensure this never happens.

Security is a critical factor in virtually all IoT use cases so systems based on Weightless-P technology benefit from leading edge data security provision. AES-128/256 encryption and authentication to the network guarantees integrity whilst temporary device identifiers offer anonymity for maximum security and privacy. OTA security key negotiation or replacement is possible whilst a future-proof cipher negotiation scheme with a minimum key length of 128 bits protects long term investment in the network integrity.

## What Weightless offers:

- Authentication to the network
- AES-128/256 encryption

.

- Radio resource management and scheduling across the overall network to ensure quality-of-service to all devices
- Support for over-the-air firmware upgrade and security key negotiation or replacement
- Fast network acquisition and frequency/time synchronisation

## COST

## Low cost and complexity

One of the most significant costs is that of the network itself and it has a number of elements. The first of these is the cost of suitable wireless spectrum. Conventionally, IoT solutions have been developed around legacy telephony based GSM technologies – GPRS, 3G and LTE. Not only are these technologies not optimal for most IoT applications but spectrum licensing costs are significant. The high cost of access to the radio spectrum inevitably translates into a correspondingly high cost of data transmission for the end user.

Weightless technology is frequency agnostic and can operate across different spectrums – importantly including licence exempt spectrum. Weightless technologies are typically deployed in sub-GHz frequencies; Weightless-N and Weightless-P operate in Industrial, Scientific and Medical (ISM) bands.

Longer range, fewer base stations

Range is another element with a strong influence on the cost of network deployment. Range translates proportionately to cell size and is one factor in the determination of the number of terminal devices that can be associated with each base station. The fewer the base stations needed in any given network, the lower the overall network cost.

Both Weightless-N and Weightless-P technologies are capable of ranges equivalent to and better than GPRS, 3G and LTE. Excellent signal propagation characteristics mean long range and excellent penetration of signal into buildings. Typically an urban range to terminals mounted internally is around 2 - 5 km whilst a line of sight rural implementation could achieve up to 30km. However, as we have seen in a previous article, range is usually not the limiting factor that determines the number of cells and base stations in a network, it’s capacity.

Excellent signal propagation in sub-GHz spectrum also enables smaller, less sophisticated and lower cost antennas to be used on endpoints while antenna location, which might give rise to incremental costs, can in many cases become less critical.

Lower CAPEX

Whether you are developing a network that you will operate or you are designing applications to connect to an existing network managed by another operator, low network installation costs lead to lower tariffs to each user of the network. The capital expenditure (CAPEX) for a network is considerably less than for an equivalent traditional cellular technology. Commercial grade basestation hardware typically costs less than USD$3000.

Lower OPEX

Weightless devices have been designed so that they do not send much data, and do not greatly load the network. Sensor-style devices may only wake up once every 15 minutes, or even less frequently. This will allow a tariff that charges only a small amount for when a device is communicating, and nothing at all when it is not, because the loading it places on the network is very small. This very low operational expenditure (OPEX) makes it possible to achieve subscription fees of a few dollars per year needed to meet typical IoT use case requirements.

Low cost hardware

Weightless specifications require no exotic, high cost components to implement at either end of the link. Commodity transceiver and microcontroller devices together with a regulator, crystal, a small number of passives and a low cost antenna mean that modules can be produced for less than USD$2 in volume. Weightless specifications have been created in such a way that they can be implemented on modules costing just USD$1 - 5 in volume.

Low cost deployment and maintenance

Fit and forget pricing of modules is complemented by the low power requirements for Weightless terminals making battery operation feasible. Battery powered devices avoid costly deployments requiring a connection to the grid whilst battery lives measured in years mean less truck roll on scheduled maintenance.

Cost of LPWAN networks is not obvious

The cost of an LPWAN network is not just the cost of a base station – in fact as we will see, far from it. For the network operator the costs per site include:

- Hardware such as the base station and antenna
- Installation, including any site engineering needed
- Site rental
- Power and telecoms to the site

There are also core network costs and the operator may also bear the terminal costs either directly or in some indirect manner. If we consider an example of a UK-wide network we can see how these costs break down.

.

Modelling a scenario

A nationwide UK network might need around 6,000 base stations to provide adequate coverage. The hardware for a base station might be around GBP£3k per base station, GBP£2k for ancillary equipment and GBP£5k for site engineering. Rental might be GBP£2k/year with backhaul and comms a further GBP£2k/year. Over a 10-year span the total cost would be GBP£50k/base station or GBP£300m. Of this only GBP£18m is the base station cost – hence the hardware cost is almost irrelevant.

If we assume 10 terminals per person that equates to about 650m terminals. At, say, GBP£5 each including packaging, etc, that is GBP£3.25bn. If the operator has to pay this cost it is clearly the dominant factor by far.

This leads to two important implications:

- Minimise the number of base stations. It is better to have a base station that is twice as expensive but with twice the capacity if this can result in fewer base stations. The additional cost at one site will be about GBP£3k but the lifetime cost saved by avoiding an extra site is GBP£50k.

- Maximise the lifetime of the terminals. If replacement can be avoided during a 10-year or longer lifespan this will provide massive savings.

Maximising base station capacity is complex and is discussed above. Another factor in reducing the number of base stations is maximising range in areas where there are fewer terminals and so capacity constraints are unlikely. This can be achieved with flexible radio systems that can adapt their modulation and coding schemes according to the received signal strength and with acknowledgements such that devices on the edge of coverage can retransmit when failures occur.

Increasing the longevity of terminals is not just about robustness but also includes:

- Long battery life – not only does this mean fewer manual interventions but also fewer times that the device is opened and contacts stressed.

- Software download capability – this ensures that if there are bugs discovered in the device software they can be addressed over the lifetime of the device. This is particularly important for security flaws which might otherwise render the device useless.

.

- Open standards – having devices conform to open standards means there is more likelihood of continued supply of network equipment and spare parts and less likelihood of a change in direction from the proprietary supplier that does not provide backward compatibility.

Weightless technologies are designed to minimise overall network costs. High capacity base stations reduce the overall network base station count and offer modes that can deliver increased range where needed. Terminals have long battery life, the ability to efficiently download a software update and the solution is the only open standard currently available for LPWAN deployment which further mitigates against cost and the risk of future OPEX increases.

## What Weightless offers:

- Low CAPEX
- Low OPEX
- Long battery life reducing truck roll
- OTA software and security upgrades ensuring longevity of endpoints in the field
- Open Standards leading to competition to sustain low costs to users
- High capacity networks reducing the number of base stations required in dense urban environments
- Long range reducing the number of base stations required in less dense rural environments

# STANDARD

## True open standard

Historically only open standards have delivered sustainable wireless technologies - for good reasons. True open standards support multiple vendors which stimulates ongoing innovation and, through greater competition, lower costs. True open standards provide for access to royalty free IP to minimise production costs. Open standards ensure interoperability between manufacturers. And robust, multiple, diverse, peer reviewed design teams from across industry leads to innovations that bring reliability and high performance to the technology at a fraction of the cost of alternatives with equivalent specification. Weightless technology has been designed on a clean slate basis, from the ground up to offer optimised performance at an unbeatable price point and avoids any legacy or backward compatibility concerns.

Why standards are important

In the world of wireless communications there are no successful proprietary technologies. Many have tried, but eventually open standards always win out. There are many reasons for this, but one of the key ones is the "two-ends" problem – every wireless link has two parts, the transmitter and the receiver, or the base station and the device. Typically, the company or person buying the base station is different from the one buying the device and neither wants to be beholden to have to buy from a particular company due to decisions made by others. Open standards allow a vibrant eco-system of suppliers for both sides of the link, enabling each party to choose their preferred supplier. There are many other good reasons for the success of standards. Standards encourage competition which results in innovative products at lower prices. By enabling a range of companies it reduces the risk of obsolescence due to a company deciding to no longer support a product. Regulators are typically more inclined to find radio spectrum for standards compared to proprietary products and the route to type approval is simpler.

Getting the technical part right

Of course, it is not enough to just have a standard, it must be technically excellent, otherwise proprietary solutions might just appear more compelling. Given that the same engineers design proprietary products as well as standards, then sound technical design should be possible. However, politics and competition can sometimes get in the way. There may be a desire for backwards compatibility which can compromise a solution. Companies in the standards body may compete to embed their own IPR in the standard even if it is not fully relevant and standards bodies find resolving conflict difficult. The result can often be a standard that has multiple modes and variants, few of which are strictly needed, but which add cost to the device. Often the best standards are developed either by small standards organisations where there are few competing manufacturers, or where consensus can be reached outside of the standards body in a mature manner and then all players work constructively around the agreed direction.

### Getting the support from industry

While small standards bodies can quickly and effectively develop excellent standards they can then suffer from insufficient industry presence with the result that there are few manufacturers of equipment. If the number of manufacturers falls too low then the benefits of multiple sources of supply might not be realised and the standard will fail to meet some of the key objectives of operators and end users. Hence, it is necessary to stimulate a vibrant eco-system with competing suppliers in all stages of the value chain. Manufacturers can tend to sit on the fence, waiting to see what their competitors do before committing to a standard whose success is unknown. Successful standards are those where a few key players are persuaded to publicly state their support and investment plans, persuading their competitors to come off of the fence. Once this process starts it rapidly becomes a virtuous spiral with each additional company that joins stimulating a number of its competitors and suppliers.

Getting those first manufacturers on board requires persuasion, leadership and a good set of contacts. It typically takes a few individuals who can evangelise the standard effectively to the decision makers in industry and who can bring together a critical mass of key players. Having the strong support of a few high-profile and widely respected individuals is key to success.

### Marketing the standard

Where standards are used by consumers they need to have a strong brand. We know to ask for “Wi-Fi” and to connect devices via “Bluetooth”. Few know that Wi-Fi is based on the IEEE802.11 standard or Bluetooth on IEEE802.15. Successful standards have a simple and memorable name and logo. Marketing of the name is more often performed by the companies making products rather than the standards body. However, the standards body needs to ensure consistency of message, simple inter-working of devices and occasionally to orchestrate concerted campaigns. Marketing is typically performed poorly by the “classic” standards bodies such as ETSI and IEEE which tend to focus more on the technical standardisation work.

### Just one standard

Standards only function effectively when there is a single standard for each application space – Bluetooth for personal connectivity, WiFi for local area networking, cellular for wide-area connectivity and so on. Where there have been standards battles such as between 4G and WiMax sales tend to suffer until the battle is resolved in favour of one standard. While competition within standards is highly desirable, competition between standards breeds uncertainty and a concern of making an investment in a standard that loses out.

.

Getting to just one standard is often problematic. There are many standards bodies that are inclined to compete and often tension between small bodies able to have a clear agenda and focus, and larger bodies able to attract multiple large companies. If a company is unable to embed its IP in one standard it may start another where it believes it can have more control. Standards are relatively easy to start and very difficult to stop. Eventually, one standard gains an edge on the others and then very quickly becomes dominant. This is because as soon as one standard appears to be winning, companies tend to coalesce around it and in doing so its success becomes self-fulfilling.

IoT and standards

Do we have the right standardisation in place for IoT? The current situation is far from a single clear standard with multiple options in licensed and unlicensed spectrum. There is fragmented support from industry with most sitting on the fence and little clear marketing to consumers as to the IoT brand. Little wonder that IoT is not currently gaining significant traction and connected devices are only a tiny percentage of predicted numbers.

## What Weightless offers:

- Carefully crafted technical standards that are not constrained by backward-compatibility nor bloated by compromises.
- Increasing support from across the eco-system.
- A clear name and branding that is memorable and appeals well to consumers.
- The marketing capability and leadership to enable the virtuous circle that inexorably leads to a single worldwide solution.
- Royalty free (FRAND-Z non-assert licensing scheme) IP minimising production costs

## Independent Research Findings

The following insights, and some of the diagrams in this report, have been reproduced with kind permission from Mobile Experts LLC from their 2016 LPWA vs LTE IoT connectivity technology landscape analysis. This report is available to purchase. A 20% discount is available to Weightless Developers. More details are available at: http://www.weightless.org/about/mobile-experts-group-lpwan-report.

- 3GPP technologies are 2-4 years away from providing a competitive solution with similar performance characteristics to LPWA technologies. The lynchpin of 3GPP strategy is the development of LTE Cat-m1 and NB-IoT technologies, both defined in 3GPP Release 13, with anticipated commercial availability in early and late 2018, respectively.

- This time-gap provides the LPWA ecosystem an opportunity to establish market presence, the success of which will be the result of a complex interplay of different factors that include foremost the ability of LPWA proponents to penetrate a fragmented market landscape with long-sales cycle.

- The LPWA ecosystem has the advantage of diversity and vitality which include startups as well as major technology players in adjacent markets that see LPWA as an opportunity to chip away at the traditional service provider market. For this reason, mobile network operators have been making investments in LPWA technologies which are essentially insurance policies on future market uptake in light of the late arrival of a standardised technology which they consider critical requirement.

- LPWANs are set to play a major role in private networks that address specific application requirements. Their success in public networks is gated to a great extent on the service value proposition and return on investment, the regulatory framework, and the competitive landscape.

- Licensed-exempt spectrum regulations strongly impact network performance and the investment required to build LPWA networks, and consequently impact the financial viability of LPWA networks. The regulatory framework in the United States is more advantageous than it is Europe where between 2x – 8x more in capital expenditure is required to achieve a similar level of service as in the US, depending on technology. The regulatory framework in many other major markets such as Japan, Korea, China, and others is still evolving.

- Our detailed analysis ranks LoRa, SigFox, RPMA, Weightless, and other standards with respect to range (link budget), capacity, and cost factors. Mobile Experts has evaluated each standard (as well as comparing them to LTE variations) using a common set of assumptions, for an unbiased fair comparison. Please contact Mobile Experts to get the complete analysis.

- The detailed analysis by Mobile Experts matches LPWA formats and LTE with 86 different applications, ranging from video surveillance and automotive cases on the high end to "smart agricul-

.

ture" and other low-cost sensors. Using our deep-dive technical analysis, Mobile Experts has been able to identify the sweet spots for LPWA systems.

We estimate that growth of LPWA devices will reach roughly 200-250 million by 2020. This growth is based on the segmentation of the market described in the full report. In the applications where LPWA has the right combination of range, power consumption, data rate and cost, we have identified a specific slice of the IoT marketplace.

Our interpretation of independent research on the projected evolution of IoT connectivity technology is that LPWAN in licence exempt spectrum and 3GPP standards will continue to coexist. LPWAN technologies in unlicensed spectrum will offer a sustainable competitive advantage in energy consumption and cost. Ultimately technologies operating in licensed spectrum will offer higher performance. The decision for developers will be made according to use case.

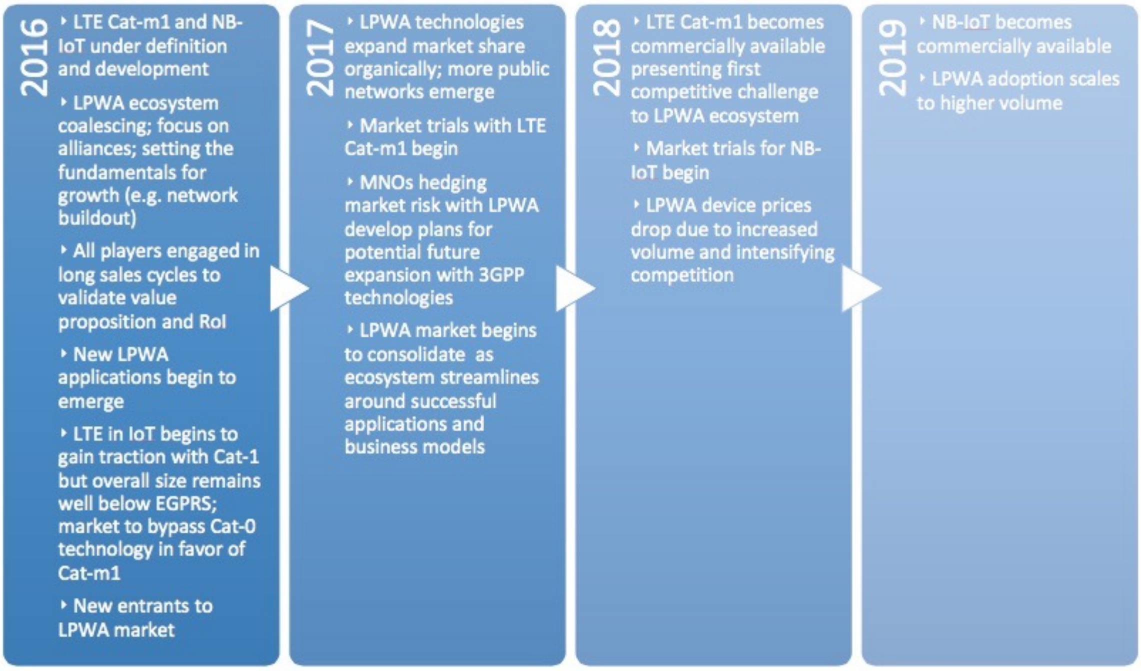

2016

- LTE Cat-m1 and NB-IoT under definition and development
- LPWA ecosystem coalescing; focus on alliances; setting the fundamentals for growth (e.g. network buildout)
- All players engaged in long sales cycles to validate value proposition and RoI
- New LPWA applications begin to emerge
- LTE in IoT begins to gain traction with Cat-1 but overall size remains well below EGPRS; market to bypass Cat-0 technology in favor of Cat-m1
- New entrants to LPWA market

2017

- LPWA technologies expand market share organically; more public networks emerge
- Market trials with LTE Cat-m1 begin
- MNOs hedging market risk with LPWA develop plans for potential future expansion with 3GPP technologies
- LPWA market begins to consolidate as ecosystem streamlines around successful applications and business models

2018

- LTE Cat-m1 becomes commercially available presenting first competitive challenge to LPWA ecosystem
- Market trials for NB-IoT begin
- LPWA device prices drop due to increased volume and intensifying competition

2019

- NB-IoT becomes commercially available
- LPWA adoption scales to higher volume

Reproduced with kind permission. © 2016 Mobile Experts LLC. All Rights Reserved.

## What’s next?

We will update this report periodically to reflect the evolution of the industry and market. We will also send you a number of emails over the next few weeks which will address each of the key technical topics you will need to consider in order to choose an IoT connectivity technology. These emails will link back to the Weightless website on which the competitive advantages of Weightless-P will be set out in detail. The 8 week, video supported, Weightless-P tutorial programme will conclude with an in-depth, interactive web seminar to answer any technical questions that remain.

The programme will also be supported by a very significant hardware offer! Details will follow.

If you have any questions at this stage you can contact info@weightless.org.

.

## 17 Critical features

We have seen above how eight fundamental parameters can define the characteristics of an IoT connectivity technology. We're now going to set out what we think are 17 critical features that you need for a successful IoT implementation using LPWAN technology. Over the coming weeks we will be setting out in technical detail how each of these is implemented in Weightless technology.

1. FDMA+TDMA in 100kHz and 12.5kHz narrow band channels offer optimal capacity for uplink-dominated traffic from a very large number of devices with moderate payload sizes

2. Operates over the whole range of license-exempt sub-GHz ISM/SRD bands for global deployment: 433/470/780/868/915/923MHz

3. Adaptive power control and data rate from 625bps to 100kbps to optimise radio resource usage depending on device link quality

4. Time-synchronised base stations and flexible channel assignment for efficient radio resource scheduling and utilisation in large-scale deployments

5. Supports paging capability, both network-originated and device-originated traffic

6. Forward Error Correction (FEC)

7. Cell re-selection (handover, roaming)

8. Fully acknowledged communications

9. Supports narrowband channels (12.5KHz) with frequency hopping for robustness to multi-path and narrowband interference

10. Supports licensed spectrum operation

11. Power consumption in idle state when stationary below 100uW

12. AES-128 authentication and encryption

13. Support for over-the-air firmware upgrade and security key negotiation or replacement

14. Brings the reliability and performance of cellular technologies at a fraction of the cost by avoiding any legacy or backward-compatibility concerns

15. Ensures interoperability between the manufacturers

16. Provides for multivendor support to stimulate ongoing innovation and minimise end user costs

17. Royalty free IP minimises production costs

.

Finally, we have set out below the Table of Contents for the Weightless-P Specification. This will enable you to appreciate the scope of Weightless-P technology and how the Standard is defined.

## Table of Contents

.

.

.

.

.