Internet of things

Published June 2015 by Rethink Technology Research Ltd

Contents

A Guide to Internet of Things Radio Protocols

We've reached a key moment in the Internet of Things (IoT). Established communications protocols are being hurriedly adapted to meet the new requirements thrown up by IoT applications, and less-known protocols are on the verge of being swallowed whole by emerging giants.

WiFi and Bluetooth are familiar faces in the consumer electronics market. With nearly every smartphone containing both protocols, they are integral parts of the majority of consumer's lives in developed economies – and are becoming increasingly important for consumers in emerging markets too.

Much less well-known are the likes of Z-Wave and ZigBee; rival mesh networking protocols that have been quietly competing in the emerging smart home markets, as well as the industrial markets – more so the case for ZigBee.

However, these protocols that are ideally suited for smart homes thanks to their low power consumption, but arrived on the scene well-ahead of any substantial demand for the smart home as a product or service – as the hardware was simply too expensive to exist as mass-market products. We are now standing at the tipping point where the smart home and its dozens or hundreds of connected devices will quickly become a mainstream offering – and both ZigBee and Z-Wave are threatened by the emergence of a new rival.

Thread is a low-power mesh network that is being promoted by Nest, the smart thermostat company that Google acquired back in January 2014 for $3.2bn. Nest went on to buy Dropcam ($555m, June 2014), for its IP security and home monitoring cameras, and Revolv's 7-radio home hub (undisclosed price, October 2014), to substantially flesh-out what is a solid foundation for a complete smart home package.

Likely to be offered as the Android smart home (Android Home, Android Wear, Android Auto, etc.), Google's interest in this market makes a lot of sense given the success of Android. Leveraging that market penetration to bring a hardware platform to market is a sensible business plan by itself – but one that is massively improved when Google's software and services are thrown into the mix too, tying more and more users more stickily into the Google ecosystem.

In conjunction with the Works With Nest program, which certifies products that will work within the current Nest product ecosystem, Thread provides a way to add connectivity to new products to ensure interoperability, using fairly common 802.15.4 physical hardware and the Thread software stack – which threatens to erode or snatch ZigBee deployments thanks to its shared MAC and PHY layers, as well as the support for native IPv6.

Consequently, the Thread Group creates another influential industry group that is pushing its own vision of the IoT roadmap. With a members list that includes a number of very influential silicon companies (including ARM, Samsung, Freescale-NXP, and Atmel) as well as significant players in the home including Yale, Tyco, and Philips, there is a lot of potential clout to bring Thread to market (by force if needs be).

But we must wait and see how Thread decides to approach the market, especially as it is not the only industry organization that is developing software or protocols for the IoT. Also on the list are the

Qualcomm-backed AllSeen Alliance and its AllJoyn device discovery and interaction framework, and a rival implementation from the Open Interconnect Consortium (OIC), which is pushing IoTivity in a very similar manner to AllSeen – to sit in between every link in the IoT connectivity chain, from lightbulb, to outlet, to home gateway.

Recent History:

A year ago, it looked like the IoT was in danger of breaking into several distinct camps, with businesses backing their chosen protocol and focusing development efforts exclusively on one. This fractured series of opposing camps would have come at the expense of the customers, who had previously been fed a vision of perfect interoperability and harmony.

Thankfully, the industry seems to have collectively moved away from this walled-garden approach, at least on the radio level, but there's still not much in the way of clarity, and is still difficult to pick winners.

What we can point towards is evidence of a significant shift that's due to occur in the coming quarters. This paper is here to guide you through the changing landscape, so that you can choose the right horse to back in this race for the IoT.

This is a tricky market to take a wide-view of, but it's something that is worth doing. I hope that your journey through these IoT protocols is less bewildering than my initial travels through the industry. We hope the following information is useful for those wishing to enter the IoT, as well as those looking to invest or formulate new business plans.

Introduction to the IoT: Why you should care

It's helpful to think of the IoT (the Internet of Things) as a verb with the following definition: to connect a previously unconnected object to the internet, either directly or indirectly, and derive some sort of value from that connection.

The Internet of Things is not a new idea or design. The benefits of connecting things to the wider internet are plain to see. What is often not so easy to distinguish is the cost and rewards for doing so.

But we've reached a point where the hardware needed to connect objects is both practical and affordable. Communications hardware has become small enough to fit inside objects that it would not have been able to 5 years ago, and the costs for adding networking capabilities to products are low enough to be seriously pursued. Just as mobile phones got small enough to fit inside a pocket rather than a briefcase, IoT devices have gotten small enough and cheap enough to be deployed en masse.

Similarly, the silicon has become suitably power-efficient to allow battery-operated units to be deployed remotely and left unattended for long periods of time. This reduces the cost of having to send a technician to change a battery or record a reading, which costs time and truck-roll expenses - which then opens up this model to new businesses or use-cases.

Why embrace IoT tech?

So now that OEMs can bring previously unconnectable objects online, they must now ask themselves why they should. This is the part of the aforementioned definition that beings the value of IoT to the table.

If we take a sports equipment manufacturer, one that makes footballs or cycling shoes for instance, there will come a point where the manufacturer can easily add a Bluetooth-enabled sensor module to a football or shoe, because the unit itself is small enough to fit inside the design and the battery-life outlasts the expected product lifespan.

With that module, which will cost a few dollars, the manufacturer can now sell a product that can actively provide feedback on the oscillation of the ball in flight, the impacts it sustains, or even the temperature of the environment it is being thrown in. With the cycling shoe, a wearer could be able to monitor the speed of their pedaling, the time and distance, as well as location if the unit was fitted with GPS capability.

Thanks to the prevalence of smartphones in developed markets, consumers are increasingly likely to desire an application that can display the data collected by their possessions. The popularity of fitness and health related wearables is proof of this fact, and the peace of mind that some consumers gain from being able to monitor their heartbeat or daily steps taken is equally likely to be created if products are created that allow a consumer to check on the status of their homes, cars, pets and families.

Anyone can see that the data generated by these modules is useful, but it is harder to convince people that this data is valuable – or rather, that a direct price can be attached to it and a viable business case written for one of these newfangled products.

The First Hop: the bones of the matter

But that's perhaps a conversation for another time. What this paper is interested in is the link between an object and the internet. Once the data is inside the IP infrastructure, it takes a well-documented trip along a series of cables to a centralized cloud installation, which pushes the bits from the shoe or ball back to an app on a smartphone.

What's more interesting for those in the IoT business, and those wishing to find out more, is the first hop from thing to receiving thing. This first hop is almost always wireless, and this is where we'll begin.

Example: The Smart Home

Looking at an example helps to understand the issue of connectivity. Seeing as the smart home and the connected lifestyle will be the largest consumer market in the IoT, it makes sense to consider how the various devices inside it will use wireless connectivity in their operations – which are rather varied.

There are a number of different criteria to consider in IoT communications. The download/uplink speed of delivery is a key one, but as you can already guess, there is a tradeoff for prioritizing bandwidth. In order to fit more bits down the pipe, so to speak, the device must use more power – i.e. it must shout louder to convey more data in a given period of time.

In order to balance the power consumption in this instance, a developer has to choose between shouting louder for a shorter period of time, or talking more quietly for a longer period. However, the much bigger factor here is whether point B is always listening for point A to tell it something, or if one or both are in an energy-saving sleep mode.

For point B to always be listening, it must always have its antenna powered on and its computational power ready to go. This requires a constant power supply, which of course will greatly reduce battery life. In order to elongate the life of a battery, point B's developer or programmer will typically instruct the device to only wake from sleep and listen a certain number of times per hour. This could be twice an hour, twice a minute, or twice a second – it all depends on what the device is.

For something that requires a very fast activation time, such as a lightbulb or door lock, a user will not put up with a high-latency or slow system for long before they resent their purchase. A light bulb needs to turn on when instructed, by an app or a voice command or even a wireless switch, and failure to do so won't be palatable.

Consequently, the lightbulb needs to be listening for a command nearly constantly, as any delay negatively impacts the user experience. This is usually not too much of a hit on performance, as the manufacturer can rely on the power supply to provide enough continuous electricity for the bulb to be always listening for an instruction. Similarly, a light switch has to be in an accessible location, so swapping a battery when needed isn't a particularly big problem.

But the big difference between the switch and the bulb is that the switch doesn't need to be listening. It can be in deep-sleep mode until it is pushed – as long as it can turn itself back on quickly enough to not inhibit performance. Since the switch has nothing else to do other than issue a command to the bulbs with which it is tasked, this should be a pretty straightforward design process.

Similarly, devices that are tasked with coordinating and controlling multiple devices in a home can rely on a suitable power supply to do so. Typically found in small standalone hubs or increasingly integrated into larger pieces of Consumer Premises Equipment (CPE, such as set top boxes or wireless gateways), the controller hub is almost always wired into a power line and connected to the outside internet via Ethernet or WiFi.

That hub might be tasked with the total control of every smart home device in the house, or perhaps just a subsystem. If users add things to it in a piecemeal fashion, it might get increasingly connected over time – but in most set ups, a central piece of equipment needs to know what's in the home network and what is doing.

Enter status updates: where the bulb tells the switch "yes I heard you, and I have now illuminated the bedroom." The switch may then pass on status update to the hub saying "I've done as you requested and now the bedroom light is on." In some architectures, the bulb could pass the message to the hub directly, and in others there could be several hops before the message reaches the coordinating device – hops that could reduce the speed of receipt.

But all the devices in the home need to provide consistent feedback to the hub. These little status updates don't require much bandwidth, but they are mission critical. Every bulb in the home needs to be able to tell the hub if it's still operational, so that the hub could tell an owner if anything needs replacing (and a notification pushed to a phone or added to a shopping list perhaps).

ZigBee chip maker GreenPeak's vision of the typical smart home of the future, with its plethora of connected devices.

Then there are the status updates that alert a device to a change, such as if a washing machine has finished its cycle, if a door has been opened, when the temperature of a room drops and the heating needs to be turned on, etc. These updates typically require action of some sort, and messages need to be passed to and from the devices on the edge of a local IoT network to the device that has been tasked with coordinating them.

The type of device will influence the speed of communication that is required. There's an obvious difference between the needs of a user to instantly change the setting of a bulb and the need to see that their laundry has been washed.

But the other major consideration for most devices is the required range of the device.

Range and Topologies: Mesh vs. Star

To return to our rather rudimentary example, a device will have to shout loudly enough to be heard from one end of a home to be heard clearly by the coordinating device – which could very well be at the other end of the building. If there are other devices in the vicinity, it is possible to pass the message from a washing machine to a security camera to a lightbulb and finally to the hub – ensuring that a device that might not be loud enough to be heard directly by the hub can still pass its message on. This network architecture is most commonly called a mesh network.

Mesh networks, as their name implies, are comprised of many different devices that are capable of passing messages up or down a chain of communication. As long as the devices are within earshot of each other, a message can reliably be passed along and received correctly at the destination. Mesh network reliability and redundancy can be increased by using repeaters or secondary hubs, so that a central hub doesn't become a single point of failure.

A well designed system will also have an element of inherent redundancy, so that mission critical systems don't fall over without central instructions – so that doors don't automatically unlock if they haven't spoken to the hub in a while, or baths don't fill themselves with more water than they can hold at maximum capacity. These two examples could prevent a burglary or flood if configured correctly, or cause the exact opposite if not enough thought is put into their design.

Critics of mesh networks often point to confirmation as a key concern. If a light bulb passes its message on to a doorbell and then returns to sleep mode immediately, there's no guarantee that the bulb's message can reach the hub intact or promptly – and potentially no way of the bulb knowing that its message is stuck in the ether and hasn't been processed.

In a star network topology, the hub sits at the center of a web with direct lines of communication to every device in the network. Devices only talk directly to the hub, so that the hub is always in control. If it issues a command that is not acknowledged, it knows it has a problem. It can try to solve that problem by shouting more loudly to the device, but it still might find that it can't reach the device it needs to talk to.

Mesh's supporters would point out that their topology would provide alternative routes from center to edge, but the mesh can introduce latency into the process – a problem that the more direct approach tends not to encounter as a command can be transmitted in one hop instead of several.

Star technologies consequently tend to have a longer straight-line range than their mesh counterparts, but often require a more powerful return-path as a tradeoff – which means that the end devices often require more power to operate. But for critical infrastructure, such as a lock or alarm, the tradeoff for range and reliability needs to be balanced with how often you need to swap the depleted batteries for fresh ones.

There are ways around this, which include several blended topologies that combine features of both, but it is worth bearing in mind that there aren't very many points on which the evangelists of either mesh or star will agree. Radio protocols are a bit like sports teams in this regard: they sometimes have rather rabid fan-bases.

So you can lengthen battery life by making the device sleep more and communicate at quieter (less intensive) levels by way of a mesh. But if the device needs a quick response that could be mission critical, you may find yourself drawn to the star despite the potential higher power consumption and potential single point of failure.

Star networks can also be simpler to diagnose and troubleshoot than a mesh network, due to the fewer potential misconnections. If the hub is operational but can't talk to a device, it know where the problem lies. In a mesh, which could have hundreds or thousands of potential pathways from point A to Z, looking for the direct cause of an outage (in terms of signal strength) can be much trickier.

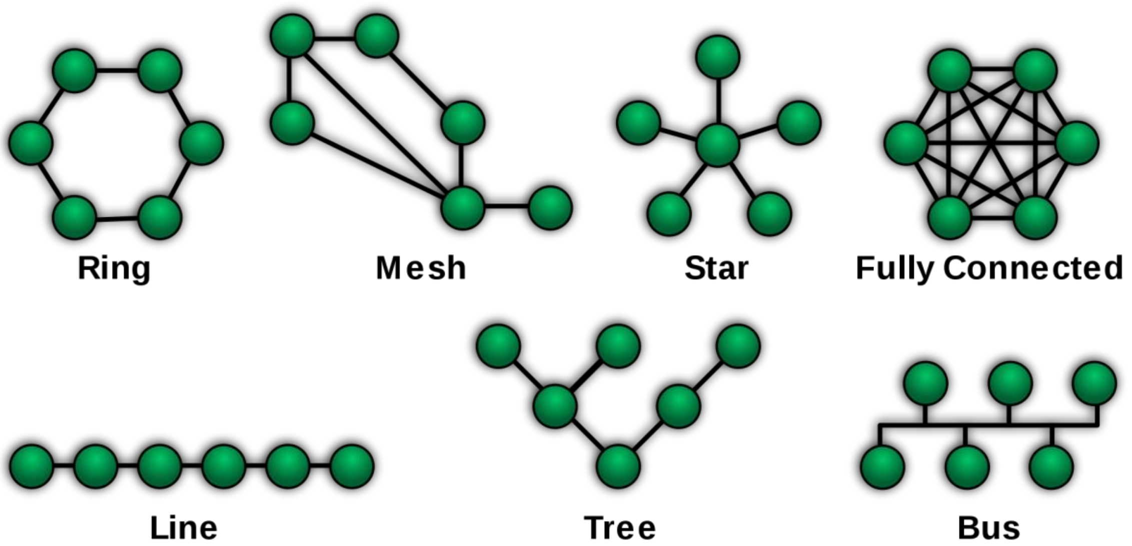

This diagram shows how mesh and fully meshed networks can follow multiple paths to a message destination, unlike star or tree networks – the two other popular network topologies in this space.

The connections themselves: onto the radios.

Those are just some of the complications that can arise when choosing the network's architecture. The radio protocols themselves have a myriad of overlapping qualities, which can make the selection process a lengthy one, and diving into the specific technical details between each one of the following protocols would take rather a long time. .

As IoT deployments are diverse in their application environments, with some only needing to cover a small home and others requiring blanket coverage of a town, no one protocol is currently suited for every use case. The following pages should help you narrow down technologies that sound most useful to your planned application, or simply act as a good overview of the IoT radio protocol environment as a whole.

So with this in mind, we here at Rethink hope the following comparisons are useful for your businesses, for choosing the basis of new project research, or building out new product lines. We'd be happy to help guide you through those decisions, and if you'd like to know more about the services we offer, please contact John@rethinkresearch.biz.

WiFi: The most ubiquitous

WiFi is found in almost every home in developed markets. You are probably very familiar with it. The wireless technology allows data to be beamed around the house, to devices several rooms apart. It has come on leaps and bounds throughout its lifetime, from its initial 1Mbps target to approaching near Gigabit speeds today – thanks to beamforming and MIMO advances, which allow multiple antennas and directional transmissions to boost bandwidth.

However, its ubiquity is also its Achilles heel, with neighboring properties and networks often encountering interference. With two neighboring access points shouting over each other, often separated by only a few inches of interior wall, WiFi networks have chalked up a patchy reputation. Parents will recognize the cries of 'the internet is broken!' and the drama that often accompanies it. While the World Wide Web hasn't fallen over, the kids are sort-of right. WiFi has a bad habit of working fine but cutting out at the most inopportune moment – often inexplicably.

This is largely why critical industrial connections are wired, or used protected spectrums instead of the unlicensed and congested 2.4GHz band. WiFi's lack of strict Quality of Service (QoS) guarantees means that it will often be trusted for baby monitors, but not carrying critical instructions for blast furnaces.

So WiFi is often seen as a 'good-enough' solution for non-critical connections, with a huge list of compatible devices and deployed CPE. You will be hard pressed to find a home nowadays that has a broadband connection but not a WiFi access point.

On the hardware side of things, WiFi modules are not the cheapest units to add to a device. While significantly more affordable than a cellular module, especially so when the cost of the monthly SIM is considered, WiFi is currently more power hungry than its rivals – but can support a much higher data rate in return.

	IEEE 802.11 – WiFi
When was it first created?	1997, but then updated in 1999 to add OFDM modulation.
Who/What invented it?	Depends on who you ask. CSIRO has been awarded $430m in damages for its patents in the 1997 implementation (despite having never sold a product), but the IEEE adopted a joint Lucent-NTT proposal as the basis of 802.11a. It's complicated; with many claiming the title.
What is the standard that defines it?	IEEE 802.11
What is the standards body / alliance / working group that governs it?	The WiFi Alliance
Is it proprietary or open source?	Open Source (free license)

What is its network topology?	Star
What is the maximum number of devices a network can support?	Many point to theoretical limit of 2007 devices per access point, but number served will be much lower, in the low triple-digits. Some access points will impose a limit due to their hardware-support for encryption keys, which can sometimes max-out at 50 devices. Multiple access points (think campus WiFi networks) can be added to solve this problem, and support thousands of concurrent users.
What is its typical and maximum bandwidth/data rate?	1Mbps - 800+ Mbps, depending on version.
What is its typical indoor range?	20m - 70m
What is its typical outdoor range?	100m – 800m+
Are all the devices IP addressable?	Yes (IPv4 and IPv6)
What frequency does it operate in?	2.4GHz and 5GHz, with some niche implementations using other spectrum bands too.
What is the frequency width?	22MHz - 160MHz
How many channels can it use?	Depends on the version, from one (22MHz), all the way up to a single 160MHz channel with 8 spatial streams in 802.11ac and channel aggregation in the 5GHz band.
What sort of modulation does it use?	OFDM and/or DSSS.
What is its maximum transmit power / irradiated power?	100mW in most cases. Some differences.
What is its peak current consumption on the device?	
What sorts of devices is it used in?	CPE, Smartphones, Homes, Businesses, Public Installations.
How many devices is it currently installed in?	Over 10 billion shipments to date, according to The WiFi Alliance.

Bluetooth Smart (Bluetooth Low Energy): oft-overlooked

Bluetooth 4.2 is broadly pretty similar to WiFi, but has less range and much less bandwidth. However, the latest version has better power efficiency and bill of materials requirements than a device that requires a WiFi module.

Bluetooth also has a similar presence in the device ecosystem, as it is installed in every smartphone – as is WiFi. Together, a combination of Bluetooth and WiFi looks like a viable platform for IoT installations, and it looks like these are the two communication protocols that Apple will be relying on for its own smart home product line, HomeKit.

But Bluetooth's reputation precedes it; infamous for flaky performance, often without explanation. Anecdotal word of mouth could be enough to see Bluetooth suffer in the consumer markets, but the Bluetooth Special Interest Group (SIG) is working to improve on the technology's reputation.

As well as focusing on creating a low-power and low-cost implementation, which the SIG bizarrely rebranded as Bluetooth Smart instead of the original and intuitive Bluetooth Low Energy, the SIG is now turning its attention to creating a mesh network implementation of Bluetooth, based on CSR's CSRmesh technology, which could open up new doors for Bluetooth in the IoT.

Recent IoT developments have largely concerned the implementation and adoption of Bluetooth beacons, devices which are fixed to an object and broadcast relevant contextual information to Bluetooth devices (usually smartphones) in their vicinity. Gimbal, a Qualcomm offshoot, and Estimote are perhaps the most prominent businesses in this domain.

Beacons are increasingly popular in the retail market as a way of increasing customer interaction and footfall, and their inventor (Apple and its iBeacons) has been forging a path with its Apple Store app and customer experience – which can allow a buyer to view information, accept promotional offers, and complete a transaction without interacting with a member of staff.

This can reduce a wage bill or free up staff, as well as reduce waiting lines in-store, but beacons are also found in sports stadiums and arenas as part of a navigation app that guides users to the nearest convenience store or restroom.

	Bluetooth Smart (a.k.a Bluetooth Low Energy)
When was it first created?	It was first merged into the main Bluetooth standard in 2010.
Who/What invented it?	Nokia, under the name Wibree in 2006.
What is the standard that defines it?	Bluetooth 4.2

What is the standards body / alliance / working group that governs it?	The Bluetooth Special Interest Group (SIG)
Is it proprietary or open source?	Open Source (Free License, GNU GPL)
What is its network topology?	Direct from device-to-device.
What is the maximum number of devices a network can support?	Currently, due to peer-to-peer implementation, it's only 2. A Master Bluetooth device can connect to 7 other devices in a Piconet, but not all devices support Piconet functionality. Work on a mesh networking version of Bluetooth has begun, with initial contributions coming from CSR and its CSRmesh technology.
What is its typical and maximum bandwidth/data rate?	1Mbps
What is its typical indoor range?	10-50m
What is its typical outdoor range?	>100m
Are all the devices IP addressable?	Yes (IPv4 and IPv6)
What frequency does it operate in?	2.4GHz
What is the frequency width?	2MHz
How many channels can it use?	Up to 40
What sort of modulation does it use?	GFSK
What is its maximum transmit power / irradiated power?	10mW
What is its peak current consumption on the device?	<15mA
What sorts of devices is it used in?	Beacons, Smartphones, Peripherals
How many devices is it currently installed in?	Upwards of 1 billion today, with 3bn per annum by 2018 according to ABI Research. Total Bluetooth shipments (including non-Smart) are estimated to have hit 10bn in March 2013, to exceed 20bn by 2016, and to exceed 30bn by 2018.

IEEE 802.15.4: father of many offspring

The IEEE 802.15.4 standard defines the physical layer (PHY) and media access control (MAC) layer for low-power and low data rate personal area networks, which have ranges that typically cover rooms instead of buildings.

It is the basis for ZigBee, which is more well-known than other 802.15.4-based technologies such as MiWi, WirelessHART and ISA100.11a. These protocols differentiate themselves in the upper layers of the OSI stack, and tend to leave the MAC and PHY alone. Interestingly, this means that an 802.15.4 radio should be software upgradeable to allow developers to switch between protocols – something that the emerging Thread standard threatens to do to ZigBee (see here).

Thread is a Google-supported approach that is itself based on 6LoWPAN – the name of both an IETF working group and its method of sending IPv6 packets over 802.15.4 radios.

802.15.4 supports both star networks and peer-to-peer mesh networks, operating in unlicensed spectrum. One of its main guiding principles is low complexity and device cost, while maintaining enough flexibility in the standard so that the typical use of 250kbps at a 10m range can be tweaked to fit the requirements – with ultra-low-power devices or very long-range devices able to be designed and meet the spec.

There are a number of lesser known 802.15.4-based protocols. ISA100's official description lists it as a "wireless system for industrial automation," MiWi is a proprietary protocol designed by Microchip for personal area networks (PANs), and WirelessHART is a wireless mesh protocol based on the Highway Addressable Remote Transducer Protocol – itself an early version of Fieldbus, another industrial automation protocol.

Currently, there isn't a unified brand that pushes 802.15.4, unlike ZigBee or Thread, which means that it is not often talked about directly in this space – despite its importance to both established and emerging protocols.

Nevertheless, it can offer a quick way for a developer or business to enter the IoT networking space, especially in niche implementations where they may want to avoid the burdens that might arise out of joining industry bodies or a standards working group. While the DIY approach isn't going to provide a new entrant with the immediate scale that adopting Bluetooth or potentially ZigBee/Thread, those other benefits might suit some companies much better than others.

	IEEE 802.15.4
When was it first created?	First defined in 2003
Who/What invented it?	The IEEE 802.15 working group.
What is the standard that defines it?	IEEE 802.15.4
What is the standards body / alliance / working group that governs it?	The IEEE.
Is it proprietary or open source?	Open Source (Free License)
What is its network topology?	Star or P2P Mesh.
What is the maximum number of devices a network can support?	Theoretical limit of 64,000 nodes per base station.
What is its typical and maximum bandwidth/data rate?	20kbps in 868MHz, 40kbps in 915MHz, and up to 250kbps in 2.4GHz.
What is its typical indoor range?	10m at 250kbps. Range can be enhanced with low noise amplifiers (LNAs)
What is its typical outdoor range?	Framework currently only covers 10m transmissions.
Are all the devices IP addressable?	Yes (IPv4 and IPv6), but doesn't support full size Ethernet frame packets, and so must be adapted to handle the larger networking packets that are larger than its 127-byte limit.
What frequency does it operate in?	868-868.6MHz (Europe), 902-928MHz (North America), 2.400-2.483.5GHz (Worldwide). These are the main ones, but 4a, 4c and 4d introduced UWB (sub GHz, 3-5GHz, 6-10GHz), 314-316 MHz, 430-434MHz and 779-787MHz (all three for China) and 950-956MHz for Japan.
How many channels can it use?	1 channel in 868MHz (EU), up to 10 channels in 915MHz (NA), up to 16 channels in 2.4GHz Worldwide.
What sort of modulation does it use?	DSSS, O-QPSK (780MHz), MPSK (780MHz), GFSK (950MHz), BPSK (950MHz)
What is its maximum transmit power / irradiated power?	100mW (in EU ISM)
What is its peak current consumption on the device?	<15mA

ZigBee: the biggest name in mesh … for now

ZigBee has undergone many revisions in its lifetime, and the current version (ZigBee 3.0) represents something of a major upheaval. In previous iterations, different profiles were released to suit different requirements, which included a separate profile for remote controls and HVAC systems.

With ZigBee 3.0, the ZigBee Alliance is pushing the mesh networking protocol as the open source (pending membership fees) rival to Z-Wave (a proprietary protocol owned by Sigma Designs, and the topic of the next chapter). However, the looming presence of google-backed Thread, which shares a common ancestor in 802.15.4 and threatens to snatch ZigBee devices with a software update, is months away from launch.

ZigBee and Z-Wave are typically positioned as rival technologies in the press. While Sigma Designs retains a license covering the production of Z-Wave chips, ZigBee currently has a number of manufacturers including GreenPeak, Atmel, Silicon Labs (Ember), Microchip, Renesas, STMicroElectronics and Texas Instruments.

Both ZigBee and Z-Wave have years on the emerging Thread, but while they may well be capable mesh networks, in the consumer space (the one that's going to get the most headlines) they seem to have arrived too early to enjoy the hype that the smart home kit is generating.

Also of note is their poor consumer visibility – as the brand is often not prominent on the packaging of the devices that they both power. Consequently, consumers are currently much more familiar with WiFi and Bluetooth, which could prove a point of contention when ZigBee and Z-Wave try and win over the consumers entering the market on the back of the rise in smart home tech.

In addition, as far as we can tell, they haven't secured contracts with the big names (Apple, Google, Microsoft, Amazon) that might push the meshes forwards in the smart home – which could see Google surge ahead with its outrageous marketing budget, or Apple resolutely stick to its current Bluetooth and WiFi strategy.

	ZigBee
When was it first created?	First standard in June 2005, with ZigBee 3.0 being finalized currently.
Who/What invented it?	Based on research concepts from the '90s, it was created by the ZigBee Alliance, an industry alliance with leading product and semiconductor manufacturers.
What is the standard that defines it?	IEEE 802.15.4 and ZigBee 3.0
What is the standards body / alliance / working group that governs it?	The ZigBee Alliance

Is it proprietary or open source?	Open standard, but requires membership of the ZigBee Alliance (which costs money), which conflicts with the GNU GPL. However, there are open source implementations of ZigBee that are available under the GNU GPL.
What is its network topology?	Star, Mesh or Tree, but most require at least one coordinator device (called a ZigBee Coordinator). The central coordinator is connected to end devices either directly or by way of routers. Strictly speaking, as end devices can't talk directly to other end devices and coordinate via the ZC, it is not a true mesh, in the truest sense of the peer-to-peer definition. However, the networks are intended to be deployed with enough coordinators and routers to ensure connectivity.
What is the maximum number of devices a network can support?	>400, but a theoretical amount as high as 64,000. You can often find frustrated developers struggling to get into the triple figures, but this could often be due to issues with the devices rather than the protocol. Typical deployments usually number 10-30 nodes, but networks with 1,000-2,000 nodes may experience latency issues.
What is its typical and maximum bandwidth/data rate?	• 20kbps in 868MHz band (EU – still not officially supported). • 40kbps in 915MHz band (Americas – still not officially supported). • 250kbps in 2.4GHz band (Global).
What is its typical indoor range?	10-20m.
What is its typical outdoor range?	• A maximum that can be measured in multiple kilometers, but these require additional hardware. • >400m a more realistic figure, but range is always a tradeoff with battery life.
Are all the devices IP addressable?	No. ZigBee 3.0 is based on ZigBee PRO, which is not IP addressable. Thread threatens that a software upgrade for ZigBee devices would make the underlying hardware IPv6 compatible, but it has not made it to market yet – and whether developers want to retroactively upgrade deployed devices is another major consideration.
What frequency does it operate in?	868-868.6MHz (Europe), 902-928MHz (North America), 2.400-2.483.5GHz (Worldwide).
What is the frequency width?	0.8MHz in EU. 26MHz in Americas. 83.MHz globally in 2.4GHz band.
How many channels can it use?	16x 5MHz bands in 2.4GHz.
What sort of modulation does it use?	DSSS, BPSK (868MHz, 915MHz), OQPSK (2.4GHz).
What is its maximum transmit power / irradiated power?	20dBm/100mW in the FCC regulations.

What is its peak current consumption on the device?	Typical sensors use 10-15mA, but some can be double that – e.g. the 31mA Freescale MC13237.
What sorts of devices is it used in?	Smart home products (including remote controls for set tops), smart meters, connected lighting, and industrial devices.
How many devices is it currently installed in?	Currently >250 million. A forecast from OnWorld predicts 1 billion units by 2020, with ZigBee Alliance Director of Strategic Marketing Ryan Maley noting its success in smart metering, with 70m UK smart meters on the books and another 104m expected to be deployed by 2020.
What does a typical message look like?	

Z-Wave: a missed chance or a key acquisition?

Z-Wave is a proprietary protocol that is currently owned by Sigma Designs, after it acquired Zen-Sys in 2008. The technology can already be found in a large number of devices in the market, but in terms of headlines, it doesn't really generate much news. It doesn't seem to be preparing itself for the imminent attack of Google and Apple in the smart home space, mentioned in the previous chapter, but it does seem to be quietly thriving in the meantime.

Nonetheless, its current install base is dwarfed by ZigBee, despite being the more established technology. Apple so far appears to be sticking with WiFi and Bluetooth for its smart home ambitions, but soon Thread will enter the mesh networking arena to compete against ZigBee and Z-Wave – which have to date been the only major players in the space. Once Nest, the current flagship device of the smart home market, commits to Thread, its partner devices and businesses will begin to encroach on the ZigBee and Z-Wave segment.

Interestingly, Samsung's new foray into the smart home SmartThings joined the Z-Wave Alliance in February 2015, and while this could just mean it wants to ensure the APIs for its washing machines will work with Z-Wave devices, Samsung's Tizen operating system has shown that the company has an unpredictable swing to it these days – which just might see it adopt Z-Wave for the platform.

However, the SmartThings hub already supports Z-Wave, so it's unlikely that Z-Wave shipments will take off on the back of a SmartThings push – itself a fairly quiet project, much like Z-Wave. This could be a very competent technology that's almost dead in the water, unless it finds a home among a business wishing to push it to the big leagues.

Z-Wave differs from ZigBee most significantly in the total number of devices that can be deployed in a single network. While Z-Wave tops out at 232, ZigBee can hit a theoretical maximum of 64,000. However, in Z-Wave's defense, we're told that functional ZigBee networks are more often found in the tens to low hundreds range, with latency becoming an issue once you hit four-figure deployments.

	Z-Wave
When was it first created?	2002
Who/What invented it?	Zen-Sys, which Sigma Designs acquired in 2008.
What is the standard that defines it?	MAC and PHY described by ITU-T G.9959
What is the standards body / alliance / working group that governs it?	The Z-Wave Alliance

Is it proprietary or open source?	Licensed by Sigma, but licensees are encouraged to participate in standards process. Alliance prefers to think of it as an open or de-facto standard. Need to purchase a development kit from Sigma, but there are some open source adaptations, which include xPL and Open-zwave.
What is its network topology?	Mesh.
What is the maximum number of devices a network can support?	232 nodes.
What is its typical and maximum bandwidth/data rate?	Began as two modes: 40kbps and 9.6kbps. But products within last two years have supported 100kbps.
What is its typical indoor range?	Certification tests provide a range of 40m, with up to 4 hops between nodes providing a higher maximum. Range will be lower with indoor attenuation and wall penetration.
What is its typical outdoor range?	Claims average range of 30m, with up to 4 hops between nodes providing a higher maximum.
Are all the devices IP addressable?	Yes, IPv6, but this requires the use of Z-Wave for IP (Z/IP) to act as a bridging technology. The communications inside the mesh will be in Z-Wave, and Sigma says this is because even the most efficient form of
What frequency does it operate in?	868.42MHz (Europe), 908.42MHz (USA), 916MHz (Israel), 919.82MHz (Hong Kong), 921.42MHz (AUS, NZ) 865.2MHz (India).
What is the frequency width?	20kHz (or 110kHz in 100kbps operation).
How many channels can it use?	Up to three.
What sort of modulation does it use?	GFSK
What is its maximum transmit power / irradiated power?	3dBm. (0dBm 1mW). Possible to achieve +6dBm in some regions with the newest Series 500 platform.
What is its peak current consumption on the device?	>1uA when sleeping.
What sorts of devices is it used in?	Smart home products mostly.
How many devices is it currently installed in?	35 million devices sold worldwide. (300 companies and 1,100 products).

DECT ULE (Ultra Low Energy): old tech with great promise

DECT will be a familiar protocol to anyone who has used a cordless phone. While it has so far not garnered much in the way of headlines, its licensed spectrum is a significant advantage over those operating in the unlicensed bands – and an upcoming 6LoWPAN implementation should make every device IP addressable.

ULE (Ultra-Low Energy) is an expansion of the DECT protocol that expands on the upper layers, tweaking the data link control, network layer, interworking unit and application layer protocol negotiation. ULE also features Quality of Service (QoS) guarantees, which are not typically found in the other low-power RF protocols.

DECT devices are also ULE-upgradeable thanks to an over the air (OTA) update, and claims that it doesn't have to conform the power requirements or time limitations that other RF protocols have to contend with. Similarly, the multiple repeaters and base stations that are common in DECT installations should help with range and latency.

The standard defines three types of devices, which are sensors (low-power consumption, long sleep times), slow actuators (mid-range time requirements, typically triggered by fixed terminals, and not requiring an immediate or very fast activation), and fast actuators (which would be found in things like light switches and alarms).

The ULE architecture aims to strike a balance between the mission-critical commands and the less time-sensitive ones. Response times can be configured from values between 20ms (which require the device to be always listening to react more quickly) and 160ms for the slow actuators. Similarly, the page cycle refresh rate can be changed, set anywhere between 160ms and 60 minutes, allowing less time-sensitive devices to consume less power by operating less frequently.

The protocol could have a future in Intel's home gateways, thanks to its recent acquisition of Lantiq, but it's early days so far for the medium.

	DECT ULE
When was it first created?	2013
Who/What invented it?	It is a fork of the DECT protocol – Digital Enhanced Cordless Telecommunications. Developed by ETSI. It was created with the cooperation of several organizations, coming from the DECT industry: DECT Forum, Dialog Semiconductor, DSP Group, Gigaset, and Vtech, in cooperation with other partners, including ETSI.
What is the standard that defines it?	ETSI TS 102 939-01.

What is the standards body / alliance / working group that governs it?	ULE Alliance is promoting the ULE technology worldwide. The transport layer standard was developed and published by ETSI; the application layer was developed and published by the ULE Alliance.
Is it proprietary or open source?	Open source. Both the specifications and the application layer implementation are open. License free.
What is its network topology?	Star.
What is the maximum number of devices a network can support?	>400
What is its typical and maximum bandwidth/data rate?	1Mbps.
What is its typical indoor range?	>50m.
What is its typical outdoor range?	>300m.
Are all the devices IP addressable?	In the current implementation of ULE, the IP termination is in the hub, while the nodes are accessed via internal ULE addressing. In the upcoming 6loWPAN version of ULE each device will be directly IP addressable.
What frequency does it operate in?	Dedicated Spectrum in: 1880-1900MHz EU, 1900-1920MHz China, 1893-1906MHz Japan, 1910-1930MHz Latin America, 1910-1920MHz Brazil, 1920-1930MHz US and Canada.
How many channels can it use?	10 in EU (1.728MHz spacing). 5 in US (1.728MHz spacing).
What sort of modulation does it use?	FDMA, TDMA, TDD.
What is its maximum transmit power / irradiated power?	Average 10mW EU, peak 250mW. Average 4mW US, peak 100mW.
What is its peak current consumption on the device?	
What sorts of devices is it used in?	Smart home products, security, healthcare, energy monitoring.
How many devices is it currently installed in?	"Currently ULE is used in a small number of devices (1 million or below), these devices were released by early adopters; however these devices do not carry the formal ULE logo – which is granted after successfully passing the certification tests. The certification program was launched in Q1'15, and by summer 2015 the certified ULE devices will hit the market." Avi Barel – Director of Business Development, ULE Alliance.

A Change of Pace: Low-Power Wide-Area Networks

The protocols we have looked at up to now have been primarily short-range, with some option to extend the device range by way of mesh-hops or using increased power to send a message to a faraway receiver. But there are two emerging protocols that are targeting long-range communications by default, and aimed squarely at IoT applications.

These two protocols are Sigfox and LoRa, and we'll dive into much more detail for each in their respective chapters. Primarily, these protocols are well-suited for long-range applications that require devices with low-power consumption. Examples of such deployments include environmental sensors such as weather stations or water quality detectors, infrastructure maintenance monitors that are installed on things like bridges, or devices like smart parking sensors.

All those applications are valuable to businesses because the sensors remove the need to send an employee to carry out the task allocated to the sensor. This also allows the sensors to provide a much larger volume of readings than a human visiting a site, at a much lower cost-per-read. In addition, this means that far more sensors can be deployed to monitor environmental conditions for less than the cost of a full-time mobile employee.

At some point, the business model will make sense for cities to replace parking wardens with sensors and cameras, or for the local authority to monitor air quality more effectively with hundreds of deployed sensors rather than dozens of larger monitoring stations. Protocols like Sigfox and LoRa aim to enable the connections between these sensors and devices with the cloud-based systems that manage them.

Sigfox: quietly confident

Sigfox is a French company that has garnered a lot of headlines in the IoT recently. With its namesake protocol software, Sigfox shares its stack with chip vendors on a royalty-free basis. It then hopes to sign up customers who use these Sigfox-compatible chips (or OTA-updated already-installed chips) to Sigfox's network – where it generates a monthly per-device subscription.

The protocol uses proprietary ultra-narrowband (UNB) technology in both the chips and the base stations, and operates in the ISM bands. Its base stations are arranged in a manner that cellular operators will find very familiar, and Sigfox says its technology is well suited for applications like automated meter reading (AMR), sensor data backhaul, and asset management.

The trade-off for its low-power consumption is found in its low data throughput. Developers have to take into account the fact that low-power technologies don't have the available bandwidth to upload images or video to the cloud, and that messages must be small in size in order to conserve power. Similarly, Sigfox does a much better job at uploading data from a device than pushing data from the cloud to the end-device.

Networks such as Sigfox (and LoRa) are well-suited for businesses or applications that have a very predictable set of needs – and enable the developer to plan their deployments to meet the specific capabilities of the available network, i.e. fixed message sizes, primarily (or exclusively) uplink traffic, non-mission-critical speed and latency requirements, and usually fixed-location assets.

If an application's requirements match those sorts of capabilities, LPWANs can provide a very affordable and quick-to-market strategy for businesses. Sigfox uses network partners to set up and operate a Sigfox network in a country (itself in France, but Arqiva in the UK, Abertis in Spain, etc.), and as a business has very little in the way of network-based overheads. Detractors say this sacrifices control over the network, but if the partner is contracted to provide a set level of coverage, you'd hope this wouldn't be too much of an issue.

	Sigfox
When was it first created?	2009
Who/What invented it?	Sigfox
What is the standard that defines it?	Sigfox UNB (Ultra Narrowband).
What is the standards body / alliance / working group that governs it?	ETSI Specifications for Low Throughput Networks (LTN). 3GPP GERAN work items.

Is it proprietary or open source?	Proprietary
What is its network topology?	Star – much like traditional cellular topology.
What is the maximum number of devices a network can support?	Claims more than 1 million messages per base station per day.
What is its typical and maximum bandwidth/data rate?	100bps EU. 600bps FCC.
What is its typical indoor range?	-142dBm sensitivity allows deep indoor reach (intended for basement-installed metering and similar applications).
What is its typical outdoor range?	Networks deployed based on average range of 30-50km, but LoS range could be much greater, with urban ranges much shorter. Sigfox says range is dependent on network deployment, rather than a single radio.
Are all the devices IP addressable?	No.
What frequency does it operate in?	Ultra Narrowband (UNB) in any ISM band. 868MHz EU, 902-928MHz US.
How many channels can it use?	Does not use fixed channel numbers. Uses Aloha access technology to allocate bands dynamically.
What sort of modulation does it use?	DB-PSK
What is its maximum transmit power / irradiated power?	25mW EU. 150mW US FCC.
What is its peak current consumption on the device?	Typically 10-45mA.
What sorts of devices is it used in?	Metering, health trackers, cold chain management, smoke and security alarms, waste management, etc.
How many devices is it currently installed in?	8 million devices on the order books so far.

Semtech's LoRa: an MNOs dream?

The clue to LoRa is in its name: Long-Range. Developed by Semtech, the proprietary approach is being pushed by the company on the back of its own radio silicon. With the LoRa Alliance, Semtech hopes to bring many more partners into the fold, and the open-source gateway protocol has already brought the likes of IBM into the mix.

LoRa is another approach to LPWAN, which was officially unveiled at Mobile World Congress 2015. While tackling the same problem as Sigfox, it operates in a markedly different manner, with different capabilities. Proponents from each side are fairly vocal on the merits of their chosen platform, and the flaws of their rival, but the technology behind both can be readily adapted to suit a wide range of application requirements.

As it stands, LoRa can be adapted to support licensed spectrum, which makes it an enticing proposition for MNOs looking to leverage frequency bands that they have exclusive rights to – while still being able to work in the ISM bands. As long as the respective national regulators are agreeable, the advantage of being able to run a low-power technology in protected spectrum are pretty clear.

Broadly, LoRa is used more in project-based deployments, focusing on a specific area or application; which is pretty much the opposite of the 'build it and they will come' philosophy of Sigfox's national coverage approach. However, Bouygues in France has recently picked LoRa for what is said to be a national project, and we'd expect other MNOs to target similar network objectives using LPWAN technology.

The chirp-based approach uses Adaptive Data Rate (ADR) tech to scale the power consumption and bandwidth to match the network conditions. Essentially, if a device is near a gateway, it can shout louder for a shorter amount of time, and act more efficiently. Similarly, far away devices will talk for longer at lower power levels in order to
speak to the gateway unit. The software-defined network will organize this in the cloud.

	LoRa
When was it first created?	2014
Who/What invented it?	Semtech, with the MAC layer developed in partnership with IBM, Actility and Microchip.
What is the standard that defines it?	LoRaWAN R1.0
What is the standards body / alliance / working group that governs it?	The LoRa Alliance – Technical Working Group

Is it proprietary or open source?	Open Source
What is its network topology?	Star – much like traditional cellular topology.
What is the maximum number of devices a network can support?	Technically no upper limit on devices, but the gateways are effectively distributing a finite amount of spectrum to the devices. As such, gateways that are supporting close-range devices are able to quickly talk to the devices and keep them off the air, enabling more devices to use the network. Gateways that have to support devices that are farther away (and require more time on the air to send their messages) will have a lower individual device capacity. Currently, there are gateways that support tens of thousands of devices each, at 300bps and at extreme ranges – according to Semtech.
What is its typical and maximum bandwidth/data rate?	0.3kbps-50kbps.Typically 1kbps.
What is its typical indoor range?	2-3km indoors, and 1-2km underground.
What is its typical outdoor range?	2-5km in dense urban areas 15km in suburbia 64km in San Francisco trial.
Are all the devices IP addressable?	IPv6 compatible mode, as well as 802.15.4g, Wireless M-Bus.
What frequency does it operate in?	Current radios can handle the entire sub-GHz band above 80MHz, so will support those apps that choose the EU and FCC sub-GHz bands.
What is the frequency width?	Typically 125kHz, of which the carrier is able to choose between 8 channels in case of interference.
How many channels can it use?	Typically 8 channels, + 2 high speed channels – although this depends on the local regulator.
What sort of modulation does it use?	LoRa's own modulation, with GFSK for the high speed channels.
What is its maximum transmit power / irradiated power?	+14dBm in most EU channels, which is enforced by ETSI. +20dBm to +30dBm in the US, enforced by the FCC.
What is its peak current consumption on the device?	+32 mA at +14dBm.
What sorts of devices is it used in?	Low-power sensors, such as parking sensors, streetlights, and gas cylinder level detectors.
How many devices is it currently installed in?	Currently, >100 integrators are offering LoRaWAN sensors. A catalog due in September 2015 should provide a pretty good estimate.
What does a typical message look like?	Typically a Preamble, then a Frame header of 12 bytes, with a message payload of 1-222 bytes, followed by the cyclic redundancy check (CRC), to check for errors in communication.

Other notable mentions

ADRadioNet:

Analog Devices (AD), a large US-based semiconductor company, is mostly known for its signal processing products. However, back in June 2014, the company launched ADRadioNet, and followed up in November with its ARM-based Wireless Sensor Network (WSN) developer kits.

The bi-directional protocol itself is justified by AD's belief that the cost of a complete standards-based solution will outstrip what a proprietary platform can offer at scale. With ADRadioNet, the smaller software stack is said to be key to lowering the BOM cost of end-devices – by upwards of 50% according to AD.

Designed for the sub-GHz and 2.4GHz ISM bands, ADRadioNet differs from other protocols by removing the requirement for routing algorithms, which require more memory to store the routing table entries required to keep the network operational.

With multiple PHY options, providing speeds from 38kbps up to 300kbps, AD says that the network range is heavily dependent on the chosen data rate – ranging from tens to hundreds of meters accordingly, with pure line-of-sight ranges for sub-GHz devices reaching multiple kilometers.

Focused around a central point (CP), the network places devices in up to 14 orbits, and can be visualized as a solar system. Messages from the outer orbits hop down the network to the center, which can then send an acknowledgement all the way out to the edge, thanks to its power supply. Extra CPs can be installed to extend the range too. A closer look at ADRadioNet can be found in RIoT here.

Link Labs

Link Labs is another player in the LPWAN game, and is pretty similar to LoRa. The company has developed PHY-agnostic software stacks that incorporate LoRa, so that it can jump ship if needed – with a long-term strategy of becoming a PaaS vendor rather than a hardware company.

Its LoRaWAN replacement is called Symphony Link, and is built to replace the MAC-layer functionality that would normally be governed by LoRaWAN. With an AWS back-end, Symphony Link lives on the gateway itself, and not in the cloud installation like LoRaWAN.

Link Labs claims its project-based network can support up to 250,000 end-points per gateway, with ranges of over 7 miles. The modules in the end devices use burst transmissions lasting less than 100ms, that draw around 50mA of current, and use an adaptive data rate very similar to LoRa. Like most of the other protocols in this paper, Symphony operates in the global ISM bands. It also supports packet acknowledgment (bidirectional communication), the usual raft of encryption measures (AES encryption, but also TLS), as well as a gateway roaming mechanism that allows up to 48 gateways to coexist in the same ISM band – according to the company.

On-Ramp Wireless

Another entrant into the LPWAN space, On-Ramp Wireless brings its Random Phase Multiple Access (RPMA) tech to the table, with gateways that it claims can cover up to 400 square-miles, catering for 16,000 end points. Pointing to its work with San Diego Gas and Electric, On-Ramp says the utility covered all the assets in the 4,500 square-mile territory using just 34 gateways. While that deployment only covers 4,000-5,000 devices, On-Ramp says it is only using 0.1% of its total network capacity.

With 172dB of link budget, On-Ramp says its approach enables underground deployments. The gateways can be multi-radio, and can support up to 80 1MHz channels. Using the 2.4GHz ISM band, and DSSS modulation, the company believes that its single-band single-radio approach provides global integration benefits.

On-Ramp provides the radio modules, gateways, and even the back-end server hardware and software if a business needs it – with On-Ramp managing the deployment in the cloud. The end-to-end approach keeps everything in-house, which should provide greater control over the deployment should anything need tweaking.

The PaaS offering has already won contracts in the utility, gas and oil markets, and On-Ramp's current go-to-market strategy sees it primarily act as the systems integrator for companies like GE – where it doesn't always interact directly with the customer.

Neul, Huawei and now NWave's Weightless-N

Neul, the Gaelic word for cloud, is the name of a promising startup that was acquired by Chinese monolith Huawei in September 2014, for around $25m. Neul developed a connectivity platform called Weightless, which targeted low-power IoT applications using the gaps in the spectrum set aside for TV broadcasts – known as white spaces, or TVWS.

After running into problems with international regulation and standardization, primarily due to the different region-specific spectrum allocations for TV found throughout the world, the Weightless Special Interest Group (SIG) turned its attention to the ISM bands, but Neul had already released its TVWS chipset, called Iceni (another nod to the company's heritage).

After its inventor was acquired by Huawei, news regarding the TVWS implantation seems to have dried up. The gaps in the TV frequencies provide a short-term solution to the much-feared 'spectrum-crunch,' with MNOs looking to the space as a way of increasing their LTE bandwidth to avoid running out of capacity.

New IoT applications would appreciate access to new spectrum, especially in the sub-GHz bands. The next few years could see something of a struggle between MNOs and the IoT – although the smart MNOs will be the ones who bolster their portfolios by embracing what the space has to offer, and leveraging their existing networking assets and expertise to do so.

After transitioning away from TVWS, towards the unlicensed ISM bands, the Weightless SIG basically turned to NWave, a UK-based wireless company, which says that the upcoming Weightless-N spec is almost entirely based on its own ISM-based protocol.

RIoT has previously spoken to NWave's CEO Jonathan Wiggin, who said that the SIG had jumped ship to ISM due to slowing progress post-Huawei. Donating its tech as a reference design for 868MHz and 902-28MHz implementations, version 1.0 of the bidirectional protocol is expected to be released in June this year.

NWave recently partnered with Plat.One to create a ready-to-go IoT networking package for enterprise customers to support applications built using the Plat.One platform. Currently, NWave claims a 5km urban range, increasing to between 20-30km in more open environments. NWave claims a 100bps data rate, and provides the radio modules and gateways to contract-based customers.

About Rethink

Rethink is expert in technology; Wireless, Video and the Internet of Things. Unlike some analyst groups, we don't sugar coat how tough the future will be. We offer consulting, advisory services, research papers, plus two weekly research services; **Wireless Watch,** a set of weekly research notes which assess the impact of events in the world of wireless and mobile technology; and **Faultline,** which provides a framework of technical understanding to help predict the future of digital entertainment and the global broadband economy, which has become required reading for anyone operating in and around quad and triple play services and digital media; and Rethink Internet of Things (RIoT), a free newsletter which covers events in the Internet of Things.

Sales contact details
John Constant, +44 (0)1794 521411
Email: john@rethinkresearch.biz
Website: www.rethinkresearch.biz