LoRaWAN®

LoRa Alliance®

LoRaWAN Empowers Very Low-power, Wireless Applications

Editor-in-Chief: John Koon

TECH IDEA
tech idea international

Version 1.0

Published by Tech Idea Research, a division of Tech Idea International, San Diego, California, USA.

TABLE OF CONTETS

1.1 Introduction of the LoRaWAN® eBook

By John Koon, Editor-in-Chief

Why LoRaWAN and LPWAN Connectivity Make Sense for the IoT

The Internet of Things (IoT) affects our lives in many ways, and to do so the "things" need a wide area network to connect to the cloud and one other. Often, many applications need only low-speed, infrequent communication. These devices can be remote, mobile, and usually battery-operated. Frequent battery replacement is both time-consuming and costly. Therefore, a solution such as a low-power wide area network (LPWAN), which can support devices for seven years or more without requiring battery changes is very attractive.

What is LoRaWAN?

The long-range wide area network (LoRaWAN) specification defines security and carrier-grade IoT LPWAN connectivity. LoRaWAN baud rates range from 0.3 kbps to 50 kbps and suffice for most remote monitoring applications. Additionally, LoRaWAN end devices are able to work with multiple networks and roam from one to another, even though these networks are run by different operators. The applications for LoRaWAN is very broad. They include utilities, tracking and logistics, smart city and parking, agriculture and farming, intelligent building, monitoring of remote things like trash, cows, water pipes and the list goes on.

Overseeing the standard is a non-profit association, the LoRa Alliance, with more than 500 members. It also manages the LoRaWAN certification program, which ensures the interoperability of all certified devices. Some of the leading member companies include Cisco, Actility, STMicroelectronics, Alibaba, Comcast MachineQ, Tencent, Cisco, Semtech, Sagemcom, Bouygues Telecom, Kerlink, Orange, Schneider, Bosch, Diehl, American Tower Brazil, and Mueller.
The official goal of the LoRa Alliance is "To support and promote global adoption of the LoRaWAN standard ensuring interoperability of all LoRaWAN certified devices".

LoRaWAN has gained strong momentum since its inception in 2015. (See *Momentum is Building for LoRaWAN*). As of beginning of 2019, the LoRa Alliance achieved a major milestone, surpassing over 100 LoRaWAN network operators globally and the network footprint continues to expand globally. As of October 2019 there are 120 operators with networks in 140 countries.
The LoRaWAN Coverage map, segmented by different network operator classifications, can be found on the home page of lora-alliance.org/

Purpose and Organization

This eBook gathers relevant LoRaWAN information into one publication for convenient access. The book includes multiple sections.

- Introduction by the Editor-in-Chief
- Interview with the LoRa Alliance Chairperson
- Interview with the Director of Certification, LoRa Alliance
- LoRaWAN Architecture Overview
- Design Articles
- Case Studies

Editor-in-Chief
Tech Idea Research
John Koon
San Diego, California
January 2020
john@techidea.com

1.2 Momentum is Building for LoRaWAN

By John Koon, Editor-in-Chief

The LoRa Alliance® holds member meetings and open house LoRaWAN Live events regularly at different locations. In the last 12 months these have included Vancouver, Canada (June 2018) , Tokyo, Japan (October 2018); San Diego, US (February 2019); Berlin, Germany (June 2019) and New Delhi, India (October 2019). Each event is well supported by members and an external audience keen to meet the ecosystem showing that the standard, now widely recognized as the de factor standard for unlicensed LPWAN, has gained a great deal of momentum and interest. The Alliance's increasingly robust ecosystem includes members from operators and silicon manufacturers, to device makers and service providers. Amazon and Intel are some of the newest high profile brands to join as members which signals strong backing from key players in the industry.

Figure 1: LoRa Alliance®member meeting and open house in Berlin, Germany showed that the technology has gained a great deal of momentum over recent years. Picture use with permission of LoRa Alliance

The open LoRaWAN® standard for IoT low-power wide area networks (LPWANs) enables smart sensors, devices and other objects to connect wirelessly via gateways that relay messages to a central network server. It is known for its bi-directional communication with end-to-end security, mobility and localization services. Perhaps most noticeable of all is its

combination of low power consumption and long range. Depending on the application, battery life can be up to 10 years and its range can extend greater than 10 miles. The applications of LoRaWAN are very broad; they cover smart cities, smart manufacturing, industrial IoT, smart homes, infrastructure, smart farming, smart metering, supply chain management, and transportation, including asset tracking. The LoRaWAN ecosystem includes providers of chipsets, modules, devices, gateways, servers (network, security, application or geolocation), networks (public, private or hybrid), cloud platforms/data management and complete solutions.

Many legacy wireless systems use frequency shifting keying (FSK) modulation as the physical layer because of its efficiency in power use. The physical layer underlying LoRaWAN is based on LoRa® Technology using the chirp spread spectrum modulation which has similar low-power characteristics but is able to communicate over a much longer distance. Chirp spread spectrum has been used for many years in military and space communication for its long-distance capability and robustness. Additionally, LoRaWAN provides bi-directional communication and supports secured multicast for applications like firmware over-the-air (FOTA) updates (Figure 2).

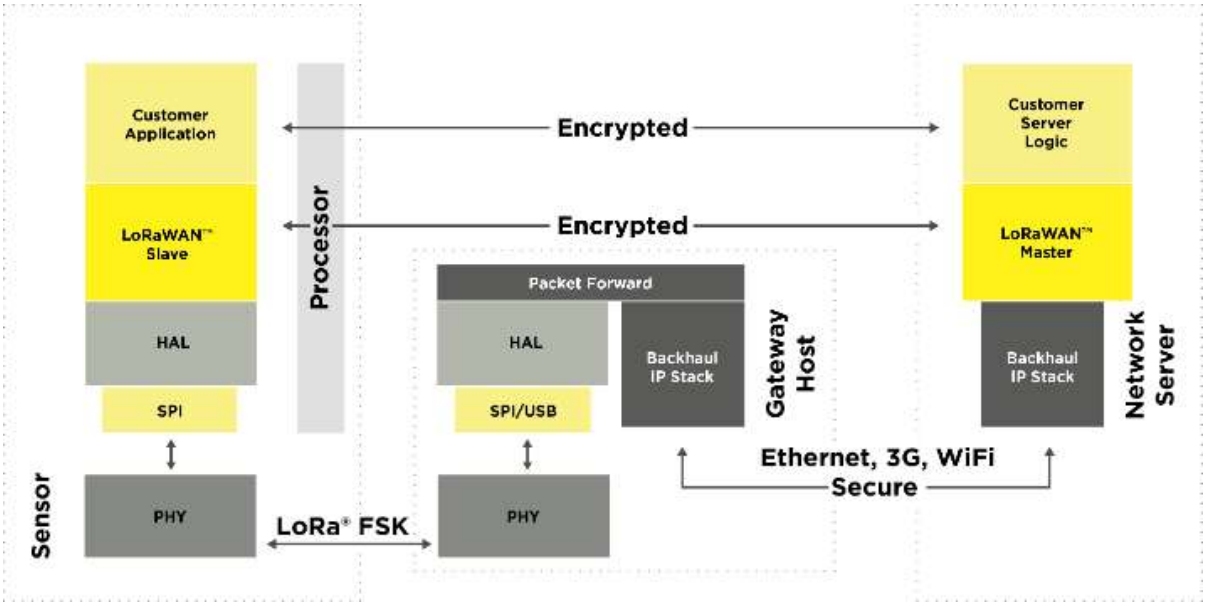

Figure 2: LoRaWAN system diagram showing how sensors are connected to the network servers. The LoRaWAN specification also provides security protocol to safeguard the network.

Behind the LoRaWAN standard is the LoRa Alliance, which is a non-profit association founded in 2015. Today it has more than 500 members and boasts millions of devices connected globally using the LoRaWAN specification. The vision of the alliance is to continue to develop and support the LoRaWAN specification and promote global adoption of the standard by ensuring interoperability for all LoRaWAN products through its certification program. Most importantly, the LoRa Alliance certification program offers multiple test houses all over the world to ensure the interoperability of LoRaWAN devices.

Momentum Is Building

In a short few years, LoRaWAN has captured the fascination of many tech companies worldwide. “Amazon and Intel recently joining the LoRa Alliance is a clear signal that LoRaWAN connectivity is gaining strong traction for IoT,” stated Donna Moore, CEO and Chairwoman of the LoRa Alliance. “All of the data generated by connected devices will enable new insights to be derived. Leading players in the industry participating in the LoRa Alliance will strengthen our efforts for end customers to realize value from this IoT data. It is my honor to welcome Amazon and Intel to the LoRa Alliance and we look forward to their support to make the value proposition for LoRaWAN even stronger.”

At the LoRaWAN Live events held in different locations during the year the members of the LoRa Alliance ecosystem showcase their products and solutions. At the Vancouver event, June 2018, a broad selection of members were exhibiting in the market place including A2A Smart City, Axatel, Blink Services, Bouygues Telecom/Objenious, Cisco, Comcast/machineQ, Digita, eleven-x, IoT Network AS, Kerlink, KPN, Loriot, M2B Communications, National Narrowband Network, Netzikon, NextField, NTT, Orange, OrbiWise, Proximus, Quaenet, Senet, SK Telecom, STC, Swisscom, The Things Network Foundation, and Unidata.
In the USA, Comcast announced machineQ, its enterprise Internet of Things (IoT) service. It has deployed LoRaWAN IoT networks in more than 15 U.S. metro areas. Most recently, in the San Francisco Bay area, where it now provides coverage in Cupertino, Fremont, Hayward, Menlo Park, Mountain View, Oakland, Palo Alto, Redwood City, San Francisco, San Jose, Santa Clara, and Sunnyvale. Sprint is said to be testing their LoRaWAN services in multiple cities. This represents a new momentum for LoRaWAN. Only a few years ago, the major operators in the USA were solely supporters of the LTE and NB-IoT solutions. The following are examples of companies who have rolled out LoRaWAN solutions.

Examples of LoRaWAN solution roll-outs

Actility launches international roaming hub

Actility, a key player in LPWAN connectivity, announced the first international hub for international IoT roaming dubbed ThingPark™ Exchange. Several LoRa Alliance members including Switzerland’s Swisscom, Netherlands’ KPN, Belgium’s Proximus, Finland’s Digita, Germany’s Netzikon and Schneider internationally connected their LoRaWAN network servers to the Actility ThingPark Exchange platform.

While roaming between different LoRaWAN operators is standardized, the process is typically expensive and inefficient, requiring operators to establish individual connections between many public or enterprise operators. Once setup on the ThingPark Exchange platform, operators will have instant connections to the other partner operators. Global roll-outs will become substantially easier, allowing companies to contract with only one operator for their IoT connectivity. The primary operator connected to ThingPark Exchange can make use of other Exchange-connected roaming partners for a complete connectivity solution for its customers.

Fully compliant with the new LoRa Alliance roaming specifications, ThingPark Exchange is a component of Actility's ThingPark Services Center, which offers a range of services to enrich and augment the capabilities of deployed networks. Actility's ThingPark multi-tenant platform is being utilized in a project with Conexxion™ for the City of Hull, in the north of England. A true "build it and they will come" Smart City infrastructure project, the company will offer paid and free services once the network is deployed.

Cypress offers silicon-based secured LoRaWAN solutions with partners

Security is a primary concern for smart city applications. The module, offered by Cypress partner Onethinx, connects to Bosch Sensortec's Cross Domain Development Kit (XDK) for Micro-Electromechanical Systems (MEMS) sensors and to the provisioning system from ESCRYPT to securely connect.

"With its long range and low power capabilities in a compact footprint, Semtech's LoRa Technology is ideal for vertical applications," said Vivek Mohan, Director of IoT in Semtech's Wireless and Sensing Products Group. "Combining LoRa Technology with Cypress' PSoC 6 MCU in a purpose-built IoT module will accelerate the delivery of leading-edge IoT solutions for smarter cities and a smarter planet."

"Devices for LoRaWAN networks need to be secure to protect the data they generate and process," said Jack Ogawa, senior director of marketing for the MCU business unit at Cypress. "We're excited to work with LoRaWAN ecosystem members to offer developers the unique combination of secure PSoC 6 MCUs and Semtech's new lower power, longer range LoRa Technology to address smart city applications."

Platform From eleven-x Enables IoT Connectivity for Unconnected Devices

Retrofitting currently deployed but unconnected IoT sensors and devices just got a bit easier thanks to a new platform from eleven-x Inc. that includes interface units that enable easy and secure wireless LoRaWAN® connectivity.

Rather than having to absorb the high costs of removing and replacing existing devices, the platform is designed to extend the capabilities of existing sensors and devices to provide real-time data collection and monitoring. The initial interface units includes those designed for use with water and gas meters (MIU-X), tank fill-level monitoring (TIU-X), and remote data collection from water level monitoring sensors (WIU-X).

The MIU-X essentially creates smart meters out of unconnected meters and the TIU-X enables cost effective remote tank fill-level monitoring, and avoids costly manual tank readings. The WIU-X eliminates the need for manual data collection without having to replace currently installed dataloggers.

Part of the eleven-x LoRaWAN wireless interface connectivity platform, the interface units are battery operated and will last for more than ten years with almost no maintenance. Installation of the interface units is a simple three step process that takes only a few minutes and does not require any wiring or specialty tools. Bringing wireless connectivity to existing assets that would otherwise need to be replaced enables substantial cost savings to utilities and other customers.

Kerlink and eleven-x Join Force to Provide LoRaWAN Tracking and Monitoring Systems

Deployed in 22 markets and powered by Kerlink Wirnet™ Stations, eleven-x has launched pilot programs in smart parking, asset tracking, water metering, fuel tank monitoring, smart waste management and other applications across Canada on its recently deployed LoRaWAN IoT network. Kerlink's stations manage millions of bi-directional messages every day from sensors and other data-collecting devices. The two companies completed rollout of the network early this year, enabling testing, development and proof-of-concept deployment of IoT solutions. Kerlink's LoRaWAN gateways are powering nationwide LoRaWAN networks in many European countries, as well as in India, New Zealand, Argentina, and Canada.

Kerlink is building a strong track record of LPWA deployments for key machine-to-machine (M2M) and IoT sectors. These include fleet management and tracking, utility metering, smart farming, smart cities and asset tracking. It also pioneered a land-based geolocation service leveraging radio and core network components that is less costly and far more energy efficient than GPS tracking.

Microchip Provides LoRaWAN Certified™ Wireless Module

For decades, wireless developers faced the dilemma of choosing between longer range and lower power consumption. LoRaWAN offers a very compelling mix of long range, low power consumption and secure data transmission, making it one of the top technology choices for battery-powered end-devices that need long range connectivity. With the growing Internet of Things, Microchip offers wireless technology solutions to address increasing demands for end-devices that need long range connectivity, low-power battery operation and low infrastructure costs for volume deployment.

Microchip's LoRaWAN Certified™ modules enable seamless connectivity to any LoRaWAN compliant network infrastructure. The RN2903A/RN2483A modules come with the LoRaWAN certified protocol stack, so customers can easily connect with the established and rapidly expanding LoRa Alliance infrastructure to create Low Power Wide Area Networks (LPWANs). This stack integration also enables the modules to be used with any microcontroller that has a UART interface, including hundreds of Microchip's MCUs. These modules feature Microchip's simple ASCII command interface for easy configuration and control, allowing for reduced development time. Additionally, the modules are also FCC and RED certified, saving significant certification costs and time (Figure 3).

Figure 3: Microchip DM164139 LoRa Mote Board is certified via the LoRa Alliance comprehensive certification program

MultiTech Offers Remote Monitoring Systems for the Energy Sector

MultiTech provides a solution in monitoring assets used in the energy sector to ensure not only the physical safety but from long-distance hackers as well. When a network in, say, nuclear energy is hacked, the damage is beyond power shut down, it could cause hazardous materials leak. Using the LoRaWAN standard for their MultiConnect devices, system operators can remotely monitor line sensors, tank monitoring, power plants, refineries and solar/wind arrays to see if the performing within the desired parameters. Additionally, the Distributed Energy Resource Management Systems (DERMS) can be used to manage load balance when multiple energy sources (solar, traditional coal, nuclear and wind turbines) are available on demand. Many of these remote sensors can operate for years without changing batteries. To facilitate the deployment of the LoRaWAN design, MultiTech also provides a LoRaWAN development kit, which includes everything a developer needs to work the LoRaWAN design. The kit includes all the needed test modules to connect to the cloud using LoRaWAN and all the cables, AC adapters with various foreign power options (Figure 4).

Figure 4: MultiTech LoRaWAN Developer's Kit comes with the test modules, cables and components.

OrbiWise's OrbiWAN™ Network Server for LoRaWAN

OrbiWise SA, a solutions provider for IoT networks, is featuring its OrbiWAN™ Network Server, a carrier-grade server capable of scaling from small local utilizations all the way up to nation-wide deployments. OrbiWAN provides the intelligence to view, run, maintain and manage the LoRaWAN network devices, gateways and interfaces, ensuring secure delivery of device data. For datacenter-based deployments incorporating millions of devices and thousands of gateways, additional servers can be added dynamically to handle increases in traffic volume and provide additional computing power. Designed for high availability and fully redundant, the OrbiWAN is managed via a web-based user interface, and provides full support for network-based geolocation of devices.

Ventia announced that it will utilize the OrbiWAN Network Server in its network. Vianet, Ventia's IoT business, provides IoT solutions including water and wastewater network monitoring, facilities management, environmental and asset monitoring, and smart city solutions.

Also of note, OrbiWise has also partnered with ESCRYPT and eleven-x Inc. to enhance data security by integrating ESCRYPT's new KMS for LoRaWAN with the OrbiWAN server on eleven-x's network. The solution addresses secure key management by storing unique keys for LoRaWAN devices in the ESCRYPT KMS. The handling of the keys ensures the inherent security of LoRaWAN is fully realized while eliminating the added expense and risk of manual key handling. Users can be assured of the mutual authentication of the device and network while the data communicated over the network is encrypted from the device through to the customer's application server, protecting the customer's application data from end-to-end.

Network provider QuaeNet Inc. announced partnerships to deliver LoRaWAN™ solutions to vertical markets.

QuaeNet is partnering with myDevices on supply chain refrigeration management for restaurants, hospitals and pharmacies. Government regulations require the monitoring and reporting of refrigeration unit statistics, thus automated temperature monitoring that provides alarm reporting and eliminates human error (or manual effort) will allow for safer deliverables. QuaeNet is also leveraging Senet's Radio Access Network (RAN) Operator services to meet the growing demand for automated tank monitoring solutions across Canada. With the extensive use of propane and fuel oil to heat residential and commercial properties, a reliable means of monitoring fuel levels of a large number of tanks is essential. By providing the LoRaWAN connectivity, QuaeNet is supporting Senet's partnership with WESROC, an industry leader in fuel tank and asset monitoring. The monitoring solutions capture tank levels throughout the day, enabling more accurate usage forecasts and more efficient and cost effective deliveries.

In a more down-to-earth application, QuaeNet is also partnering with Sensoterra, maker of soil moisture sensors. QuaeNet will be offering the solution that delivers real-time data regarding moisture levels in soil. Sensoterra's sensors help growers make more informed irrigation decisions, enabling them to provide the right amount of water at the right time, conserve water,

and even improve crop yields. The remote sensors could also help to reduce water waste in public landscaping and golf course maintenance, as well as consumer usage.

Semtech and Cypress Semiconductor collaborated on smart city solutions

Long range, low power, and security are the essential requirements of a Smart City LoRaWAN solution. To that end Semtech and Cypress Semiconductor collaborated on a compact, two-chip LoRaWAN based module. Deployed by Onethinx, the highly-integrated module is ideal for smart city applications that integrate multiple sensors and need to operate in harsh radio environments.
Using Cypress' PSoC® 6 microcontroller's (MCU) hardware-based Secure Element functionality and Semtech's LoRa® devices and wireless radio frequency technology (LoRa Technology), the solution enables a multi-layer security architecture that isolates trust anchors for highly protected device-to-cloud connectivity.

The Onethinx module utilizes the integrated Secure Element functionality in the PSoC 6 MCU to give each LoRaWAN-based device a secret identity to securely boot, on-board, and deliver data to the Cloud application. Using its mutual authentication capabilities, the PSoC 6 MCU-based LoRa-equipped device can also receive authenticated over-the-air firmware updates. Key provisioning and management services are provided by IoT security provider and member of the Bosch group, ESCRYPT, for a complete end-to-end, secure LoRaWAN solution.

The original article was published by Embedded Computing Design in 2018. Updated in 2019 by the editor.

INTERVIEW

2.1 One-on-one Interview with the Chairwoman of LoRa Alliance

By John Koon, Editor-in-Chief

Donna Moore is the CEO and Chairwoman of the LoRa Alliance. She has nearly two decades of experience in launching new companies and growing businesses across a variety of industries and competitive environments. Before joining the LoRa Alliance, she was CEO of SpireSpark, a company that designs, builds and manages worldwide certification, compliance and conformance programs. Prior to that, she served as the executive director of the Digital Living Network Alliance (DLNA), where she successfully led the global market adoption of DLNA with 4 billion certified devices deployed worldwide, established strategic alliances across industries, doubled Alliance membership, and increased member participation. She holds a Bachelor of Science degree from San Diego State University.

1. **What is the mission of LoRa Alliance?**

   The LoRa Alliance has a vision to support and promote global adoption of the LoRaWAN standard ensuring interoperability of all LoRaWAN Certified^{CM} devices. The LoRa Alliance strives to bring together its members to closely collaborate and share experiences to promote and drive the success of the LoRaWAN protocol as the leading open global standard for secure, carrier-grade IoT low power, wide area networking (LPWAN) connectivity. With the technical flexibility to address a broad range of IoT applications, both static and mobile, and a certification program to guarantee interoperability, LoRaWAN has already been deployed by major mobile network operators globally, with continuing expansion. Our members come from organizations of all types around the world addressing all aspects of the ecosystem. Members include multi-national telecommunication companies, equipment manufacturers, system integrators, sensor manufacturers, entrepreneurial start-ups and semiconductor companies. In the Americas, APAC and EMEA, our members develop, deploy and use the technology across countries and continents, driving the implementation of the IoT.

2. **Can you provide a short history on why the alliance was formed?**

   The LoRa Alliance was formed because a small group of companies came together in 2015 united in a belief that the time of the Internet of Things is now. To make it a reality, these companies recognized that standardization and a strong, growing ecosystem would be the only way to drive volume deployments for the low power wide area networks projected to connect 70%+ of the predicted IoT volumes. The companies united around a goal to standardize LPWA networks with the LoRaWAN specification and set out to create a certification and compliance program to ensure device interoperability. This approach means that LoRaWAN end-devices are able to be deployed in multiple networks and roam from one network to another irrespective of network infrastructure or operator. Since 2015, the LoRa Alliance has become the fastest growing technology alliance; we already have more than 500 members since launching our operations. Members range from technology leaders to leading product companies and many SME's and startup companies all adding significant value to the fast growing LoRaWAN ecosystem. Our members also include the largest mobile network operators who are deploying public networks using the LoRaWAN standard.

3. **I have seen the LoRaWAN (Long Range and Wide Area Network) technologies have been growing at a rapid pace, what are the driving forces behind it?**

   There are many market drivers for the LPWAN space, in fact, what we're seeing is that the market is really unlimited. There are so many potential IoT applications, I feel strongly that we're just at the tip of the iceberg. What's interesting is that even within a use case, different technologies can be deployed to meet different market requirements. We also continue to see new applications for LoRaWAN where initially people might not have thought that LPWANs generally or LoRaWAN specifically could be effective – think of those that require firmware updates over the air or geolocation – both technologies have been designed into the specification and are being used in deployments. Ultimately, the fact that companies and cities can achieve a solid return on investment is a major driver. Using LoRaWAN for enterprise applications has taken off significantly in the past year – largely because it can do the job that companies need, as well as offer the ability to own and operate one's own network. Finally, the fact that many of the 'big players' are engaging with the LoRa Alliance and actively developing solutions around LoRaWAN is helping propel the specification forward.

4. **Can you give us a sense of how big the LoRaWAN market will be in 5 years?**

   Market projections for the next 5 years and beyond and exceptionally strong. Lee Ratcliff, Senior Principal Analyst, at iHS Markit presented at our LoRaWAN Live event in Berlin [June 2019] and stated that "As of this year, LoRaWAN is by far the leader in LPWAN". From his market report he predicted that by 2023 LoRa/LoRaWAN connected end devices would be >700m and that 86% of all LPWAN end devices would be down to LoRaWAN and NB-IoT.

LoRaWAN holds a strong position now and into the future because members and end customers have already moved to mass deployment where other technologies are still at POC stage. The success to date is backed by examples of strong ROI and as the number of applications continues to grow, the growth will scale proportionally.

According to ABI Research [Whitepaper published 2019: LoRaWAN®, Competitors or Complementary], LoRaWAN has already proved its massive adoption in the following verticals.

- Utilities connecting smart meters for gas and water utilities.
- Smart buildings for environmental monitoring and occupancy knowledge.
- Logistics / Asset tracking for visibility and traceability of assets across a larger portion of the supply chain that extends from indoor environments to yard environments, and even across metropolitan areas and regions using a single technology.
- Industrial and smart manufacturing for improving visibility on production flow, monitor machine health to reduce downtime, view asset utilization, and study overall operational efficiency.
- Smart agriculture for monitoring soil moisture or livestock condition to improve crop yield or dairy yield; and the creation of affordable WAN networks to collect sensor data in place of cellular networks that may not be available

LoRaWAN is regarded as the de facto unlicensed LPWAN of choice for those applications that truly need low power, a long lifetime and don't need to transmit large volumes of data.

5. **When should LoRaWAN be used? Can you give us a use case that you would consider to be a perfect fit for LoRaWAN?**

LoRaWAN really is a very flexible standard in that it affords choice of deployment model. IoT solutions providers and end customers can utilize either existing public operator networks or install a private network which means that there isn't a reliance on existing network infrastructure. In addition, end users have options on business models either capex or service models and easier installation with existing legacy systems,
As I mentioned previously, the enterprise use case is very strong for LoRaWAN, both because it does the job that it needs to do, but gives companies and plant owners the ability to own and maintain their own networks – and, perhaps more importantly, their own data. We are seeing significant growth in private enterprise networks triggered by a huge volume of applications where LoRaWAN is a perfect fit – tracking, logistics, metering, really the list is endless. I should note that smart agriculture is another key use case where LoRaWAN network model flexibility is proving to be a key differentiator – when you need to receive data from remote areas of a remote location, no technology is better suited to the task.

6. **There are comparable offerings on the market such as Narrow-band IoT, LTE, SigFox, Wi-Sun and others. Is one better than the other?**

To an outside observer they may appear comparable, but each really has its unique points of differentiation. And fundamentally, the choice of technology comes down to the application requirements and business models to achieve the necessary ROI.
It is clear that in order to deliver massive IoT, within a broader 5G ecosystem of the future, there will be no single technology to cover all applications. This opens the door for a much broader ecosystem of “best fit” technologies to support the all applications. LoRaWAN already has the greatest variety and quantity of sensors and end to end solutions available today, with deployments in more than 140 countries worldwide. It is also essential to be future-proof and backed by a strong and varied ecosystem ensuring constant development and innovation as required.
Beyond availability, however, LoRaWAN has been massively adopted as a ‘LPWAN standard by design’ – due to efficient battery consumption, cost efficient roll outs enabling diverse business models like hybrid private-public networks. In comparison, 5G new radio (5G NR) priorities to date has focused on broadband services and critical communications and broadband services. The current 4G LTE IoT technologies Narrowband-IoT (NB-IoT) and LTE-M are also the LPWAN IoT technologies of the 5G ecosystem for some years to come, and will continue to be complementary to LoRaWAN as they are today.

Several Operators using unlicensed spectrum such as Multi Services Operators (cable companies, TV broadcasters, fixed operators, Internet Service Providers) have also adopted LoRaWAN as their legacy massive IoT technology. Unlicensed ISM-bands empowers these operators to leverage the agility of LoRaWAN as a cost efficient LPWAN technology meeting their performance requirements.

Private deployments are another reason why LoRaWAN is complementary to cellular IoT. LoRaWAN being an option of choice for enterprise customers and cities needing flexibility and agility to meet their specific business needs. The standard is well suited for cost efficient indoor deployments, penetrates through concrete, offering a very compelling business case in the ‘last mile’ scenario, often incorporating cellular for reliable low touch data backhaul.

Several leading Mobile Operators across regions already use LoRaWAN and cellular IoT as complementary LPWAN technologies to serve different customer business cases. LoRaWAN may be rolled-out on a water meter, smart city or smart building project where cellular IoT can be deployed for an electricity smart grid project requiring high throughput and frequent data transmission to feed analytics into the cloud.

LoRaWAN, supported by members of the LoRa Alliance and the other partnering Alliances, already seamlessly interconnects with cellular IoT at the data management level (application layer) and will continue to look openly at the best options to interconnect and collaborate with 5G ecosystem

What is key is having a strong certification program to ensure interoperability and having

a large ecosystem that offers enough options to meet market demand and ensures future proofing of a long-standing specification. The LoRa Alliance continues to invest heavily in providing a robust certification program and has recently launched the LCTT [LoRaWAN Certified Test Tool] to further support device makers to pre-test their devices before they move forward to formal certification. This is a key benefit afforded to members who join the LoRa Alliance.

7. **The LoRa Alliance members are growing. Recently Amazon and Intel have become members. If a company is still considering whether to join LoRa Alliance, what would you say to them?**

   Don't get left behind! With the IoT driving fast, LoRaWAN is experiencing a depth and breadth of deployments, this is the time to engage and join the LoRa Alliance. One of the benefits of LoRaWAN being an open standard is that its development and maintenance is member-driven. This means our members determine the future of the specification and when and how it develops – so joining gets a company a seat at the table in this very exciting market. We take the LoRaWAN Certified^{CM} mark very seriously, as it means that the product has been developed according to our specification and will be interoperable with other devices, and perhaps most importantly meet the requirements for IoT devices. The worst thing in the IoT space would be to have to dig up malfunctioning devices or to replace batteries every six months, so there is tremendous value to our members who offer LoRaWAN Certified products and can market its value. Finally, we're serious about networking! Feedback from our members is consistently positive about the business opportunities they've secured by joining the Alliance, and participating in our Member Meetings and other events. We have a huge ecosystem of vendors developing products and solutions across the value chain, and joining us offers an easy way to get in front of prospective customers and suppliers alike. A key benefit of becoming a member of the LoRa Alliance is the opportunity to develop strategic partnerships, collaborate and network to deploy successful solutions.
   Our Marketing Committee have launched the "LoRaWAN Showcase" in the last month which is the official online catalog of all LoRaWAN Certified^{CM} devices and other products and services provided by members of the LoRa Alliance. This will both aid the end customers looking for the right certified products and partners from our ecosystem for their solutions and also provides valuable promotion for our members to connect with new business leads to deliver more solutions to the market.

LoRaWAN CERTIFICATION

3.1 One-on-one Interview with the Director of LoRaWAN Certification

By John Koon, Editor-in-Chief

Derek Hunt is a Director of Certification for LoRa Alliance and Certification Committee Chair. He has been working on LoRaWAN Systems and specifications since 2014, at Semtech and Actility prior to become part of the LoRa Alliance.

Prior to this Derek has had over 30 years of telecommunications experience working in different Hardware, Firmware, Software and ASIC development roles and a variety different management roles. Theses have included the management of large multi discipline engineering teams, Mobile Network rollout groups, as well as change management and outsourcing management. Derek has also had several consultant roles assisting mobile phone operators such as Vodafone, Orange and T-Mobile during this period, working for Ericsson, Marconi and GEC-Plessey Telecommunication (GPT) at locations in Europe and North America.

1. **The purpose of the LoRaWAN certification is to ensure device interoperability, can you describe how the process work?**

   The end-devices designed using the LoRaWAN® standard are able to work with multiple networks and roam from one to another even when these networks are controlled by different operators. To ensure that each device will operate within the network and with each other flawlessly, it is important for a LoRaWAN device to go through the certification program managed by the LoRa Alliance to ensure interoperability. This process includes a comprehensive test procedure and covers many aspects including both the functional and MAC layer tests. The steps include:

   - Contact a LoRa Alliance Authorized Test House (ATH) for a quote.
   - Complete the certification question available on the LoRa Alliance website Members Area or from the ATH.
   - Prepare your product(s) for certification.

- Products must fulfill the latest LoRaWAN Specification and Regional Parameters Document.
- Products must fulfill relevant regional LoRa Alliance End Device Certification Requirement Document (see suite of regional tests listed above).
- Deliver your product to the ATH. Device should be ready for Over the Air activation or already personalized.
- The ATH will perform the certification tests and provide you and/or the Alliance with the results.
- The "Pass" test results are provided to you and/or the LoRa Alliance.
- LoRa Alliance will review the test results and issue a certificate of LoRa Alliance Certification.
- Results and basic product information are released on the LoRa Alliance website. The release date can be customized to align with a product launch date, if requested. Data from the questionnaire that should not appear on the website should be marked as confidential.

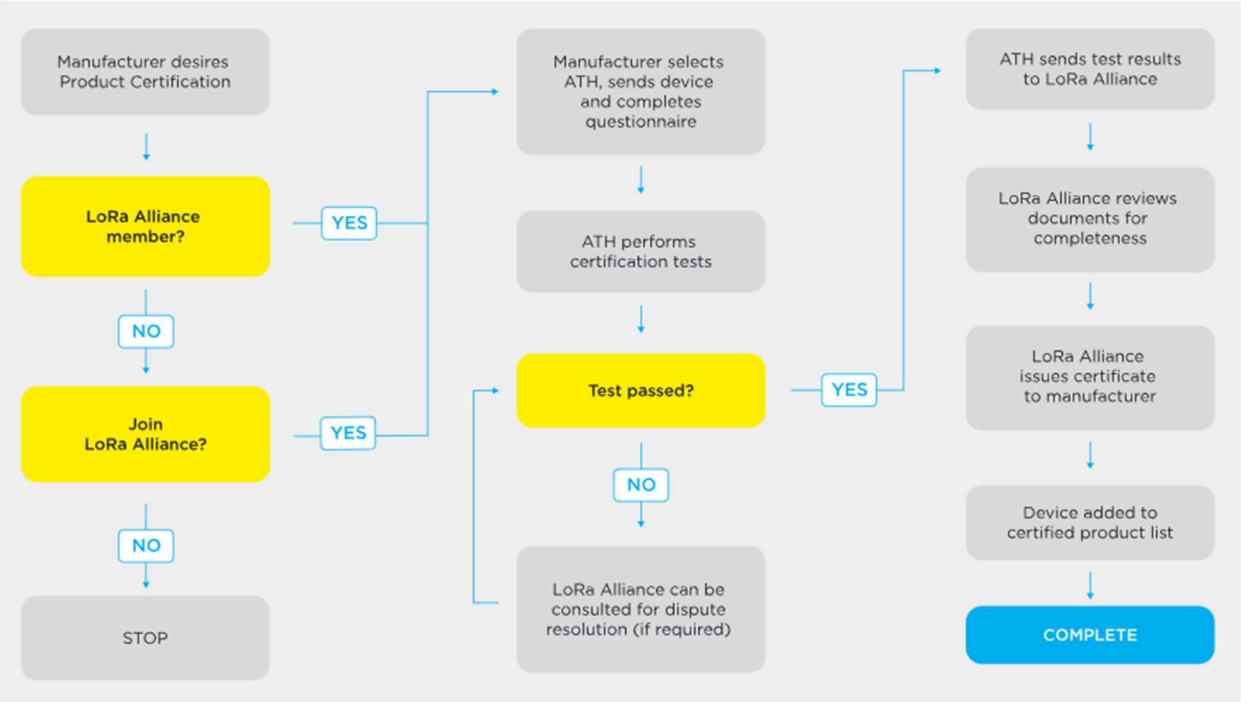

Figure 1: LoRaWAN® end-device certification process

Additionally, as each device must operate within the constraints of the ISM band for each region the device will be tested according to ISM Band regulatory requirement and will use frequencies ranging from 400MHz to 928MHz Full details of the operating frequencies in each region can be found in the LoRaWAN Regional Parameters specification produce by the LoRa Alliance.

2. **Besides guaranteeing interoperability, what other benefits does the certification provide?**

   There are additional benefits of marketing and promotional assistance. They include product promotion in LoRa Alliance collateral and highlight in LoRa Alliance demonstrations at tradeshows, listing on the LoRa Alliance website, social media and Alliance newsletter and the right to use the official LoRaWAN Certified^{CM} logo. The logo is a mark of quality.

3. **Who will be doing the certification?**

   Multiple LoRa Alliance authorized test houses are located in multiple countries for the convenience of the members. They include USA, Spain, Finland, Germany, the Netherlands, Japan, Taiwan, and S. Korea. Additional test houses will be added in the future.

4. **How long does it take to perform a certification test?**

   The test houses are very efficient and the typical turnaround time is one week.

5. **If a certified product has gone through a product update such as firmware change, what steps would be involved to be re-certified?**

   If the change is only at the application level without affecting the LoRaWAN protocol, then you would need to prove a case of certification by similarity with one of the following cases for Certification-by-Similarity:

   - Case 1 - Module Integration: the variant device to be certified embeds a LoRaWAN Certified^{CM} module.
   - Case 2 - Module Family: the variant device to be certified is a module from the same family of LoRaWAN Certified^{CM} modules.
   - Case 3 - Device Certification-by-Similarity: the variant device to be certified uses the same module as another LoRaWAN Certified^{CM} device.

6. **Any updates on the certification program?**

   Yes, to improve the overall efficiency of certification, LoRa Alliance has introduced an enhanced certification program for device manufacturers, network operators and end customers. This single certification process covers conformance, interoperability and RF testing.

   Additionally, a new LoRaWAN Certification Test Tool (LCTT) is now available. It allows the device manufacturer to pretest the full testing and regression testing before submitting to the test house saving additional time.

   At our most recent event in New Delhi, October 2019, we announced two new Test Houses opening in India with TUV Rheinland and Dekra now providing local test services.

LoRaWAN SECURITY

4.1 LoRaWAN Security Basic Part 1: Server and end device relationship in LoRaWAN network

By John W. Koon, Editor-in-Chief

Introduction

Providing end-to-end security is the ultimate security goal of every wireless network. Each reported successful cyberattack serves as a reminder that achieving security is easier said than done. Let's first examine the process. Most of us know about using strong passwords and encryption. The most vulnerable security component is key management during device activation and authentication. As with physical properties, if you hold the key to the front door, you can get in. The same concept applies to IoT applications. It's vital to manage the keys to ensure no intruders have access to your key copying or rekeying it. A smart city may have 20,000 smart streetlights (end devices) that need to be connected to the server (controller) for the first time. And connecting 20,000 end devices through a process called device activation and authentication is no trivial task. Simply put, authentication is a process to ensure "*You are who you say you are.*" A device can easily present an ID that looks real to the server. But is it? If a fake ID is accepted, the whole network is compromised and hackers can steal data information without the knowledge of the end device owners.

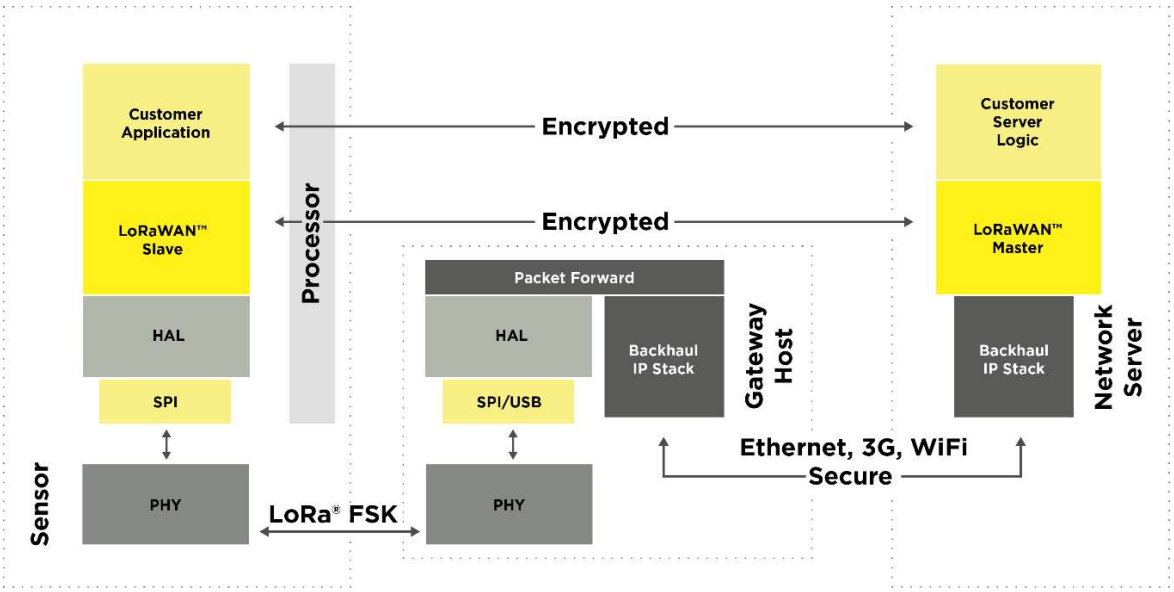

Figure 1: The LoRaWAN block diagram shows how end devices (left) connect to the network servers (right) via gateway hosts (middle). (Image courtesy of LoRa Alliance)

How LoRaWAN network provides end-to-end security

A *long-range wide area network,* commonly known as LoRaWAN, is an open, standard-based and fast-growing network platform. In just a few years, it has grown from a few network providers to 100 worldwide. Like most long-range networks such as SigFox and NB-IoT, LoRaWAN connects devices and gateways to servers. It is capable of bi-directional communication with end-to-end security, while some applications or networks only offer unidirectional communication. Typically, these low-power networks (LPWA) including SigFox, NB-IoT, LoRaWAN and others support devices with 5-10 years of battery life, making LPWA solutions very attractive to developers. The broad-based applications include smart homes, smart grids, infrastructure, smart farming, industrial IoT, smart cities and smart manufacturing (also known as Industry 4.0).

As shown in Figure 1, the blocks on the left represent end devices, which can be sensors or edge devices. In the middle are blocks representing a gateway and on the right the blocks represent customer owned application servers or network servers owned by network providers. These servers may be located remotely. Other terms used by LoRaWAN include application server, network server and join server.

- An application server hosts the application software, which controls or communicates with the end devices.

- A network server is owned and used by a network provider to forward data traffic between the application server and the end devices.

- A join server is an independent server owned by a third party other than the network provider. (Sometimes, the join server can be an independent server owned by the same network service provider.) This is used to keep the root key. Some users prefer this method to increase network security. With the key kept by the join server, even though the network provider's platform is hacked, the user's information is secure because the hacker does not have access to the key kept outside the network provider's domain.

In part 2 of the series, we will go over how authentication work to achieve end-to-end network security between the end devices and the application servers.

Original article was published by WTWH Media

4.2 LoRaWAN Security Basic Part 2: Authentication and "Key" Management in LoRaWAN

By John W. Koon, Editor-in-Chief

Device authentication and key management

Each LoRaWAN device is given a unique 128-bit AES key (called AppKey) and a globally unique identifier (EUI-64-based DevEUI) for authentication purposes. For LoRaWAN, there are two ways for an end device to initiate a connection request to the application server: over-the-air activation (OTAA) and activation by personalization (ABP). Successful connection involves authentication using multiple keys after acceptance of the activation request. Only after the activation process is completed can encryption be applied. Otherwise, either the end devices or the application server will not have the key to decrypt a message. It's as if two countries engage in a battle, and the commander-in-chief delivers a secret message to the field operator. Not wanting the enemies to know the content, the message is encrypted. But the field operator must first have access to the method of deciphering (the key) the secret message. Otherwise, the message can never be deciphered. But if the method of deciphering (the key) is stolen by the enemy, the secret message is not a secret anymore.

What is a key? It is a series of long codes. When it used properly, both the application server and the end devices can decipher the encrypted messages both ways. Note that LoRaWAN specification 1.0x specifies one root key is issued while specification 1.1 specifies two root keys are issued. But the concept of key management remains the same.

How do OTAA and ABP work?

In OTAA, multiple keys are involved. Here is how it works.

Step 1: The end device sends an unencrypted activation request known as the *Join Request* to the application server. Usually it will go through a gateway device and the network server before reaching the join server. It is similar to sending an email. The email will go through your router, then the modem and through the servers of your network provider such as ATT, Spectrum Cable or others. The join request is a long code contains the device information, DevNonce (a unique random number) and a message integrity code (MIC). DevNonce, the unique 16-bit random number generated by the end device, can be used only once to ensure that the hacker cannot use it again to prepare for the next attack. The MIC is created using an algorithm to create a checksum. This algorithm uses the 128-bit CMAC-AES encryption with the AppKey (issued to the end device). Note that at this point, no encrypted data has been generated.

Step 2: The gateway device, upon receipt of the Join Request, will send back a Join Accept to the end device. This Join Accept message consists of an AppNonce, information generated by the application server, an ID and other relevant information. Note that that the Join Accept message is encrypted with the AppKey as described in Step 1 above.

Step 3: Upon receipt of the Join Accept, the end device can now derive a new key called a "session key." This key is created with random Nonces (the DevNonce and the AppNonce) information received back from the Join Accept. Note that the "session key" is actually a set of keys called NwkSKey and AppSKey. At this point, the encrypted data communication can be initiated.

Compared with OTAA, ABP is static and much simpler. Each device and server (network and application) is issued a set of keys one time. So long as the data messages are encrypted and decrypted using the same keys, the data transfer will be successful. The potential problem of ABP is that those keys are static. OTAA provides stronger security by means of using ephemeral session keys.

Summary

The open, standard-based, *long-range wide area network* commonly known as LoRaWAN provides low-power connections for many applications including smart homes, smart grids, infrastructure, smart farming, industrial IoT, smart cities and smart manufacturing. Each end device or edge device will need to go through a connection and authentication process to establish trust so the application servers will know "*You are who you say you are.*" There are two different methods: OTAA and ABP. OTAA has additional overhead but more secure. ABP is static and simpler, but the downside is that once the keys are stolen, the device is compromised. By comparison, OTAA uses "session keys" in which each communication session will make random generated keys, and hackers cannot reuse the stolen keys. Questions remain, however. What is preventing hackers from stealing information or pretending to be the real device during the authentication process. How secure is the action and authentication process? In upcoming articles, we will explain the pitfalls of the network security and explore the hardware and partition approach to enhance network cybersecurity.

Original article was published by WTWH Media

4.3 LoRaWAN Scalable & Secured Device Activation

By Raphael Apfeldorfer, Director Advanced Technology, Actility

LoRaWAN enables device activation workflows that are more robust and scalable than most other IoT technologies. Power efficient procedures that do not require prior peering enable direct activation of the devices in the field once optimized out-of-band provisioning processes have been put in place. Security is built into LoRaWAN specifications, which restrict key distribution on a need-to-know basis and enforce security by using standard secure hardware elements.

LoRaWAN was designed from the beginning to accommodate Low Power Wide Area (LPWA) networks and low-power, low-cost devices. Based on these requirements, device activation on the network is designed to minimize the amount of data sent over the radio layer. Two activation modes are defined in the LoRaWAN standard, Activation by Personalization (ABP) and Over the Air Activation (OTAA). In the first mode (ABP), activation is done in two steps: device personalization and device provisioning. The activation information is directly personalized and provisioned for a particular network, but this is not recommended for scalability unless it's for specific use cases like devices supporting only uplink communication. In the case of OTAA, device activation is negotiated over the radio. The third step, activation, allows more independence from the network during the personalization step and high scalability, as we will discuss in this article. This article will only cover OTA Activation, please refer to the LoRaWAN specification for more details about ABP.

How are LoRaWAN End-Devices Personalized?

LoRaWAN end-devices security is handled during device personalization. This step is frequently handled during manufacturing and does not require end customer actions like inserting a SIM card into the device. This allows low-cost and efficient device deployments as any manual actions are very costly in IoT. Personalization is done by configuring three key parameters: DevEUI, JoinEUI (previously called AppEUI in LoRaWAN1.0 and renamed in LoRaWAN1.1) and root keys (AppKey and optionally GenAppKey in LoRaWAN1.0, AppKey and NwkKey in LoRaWAN1.1).

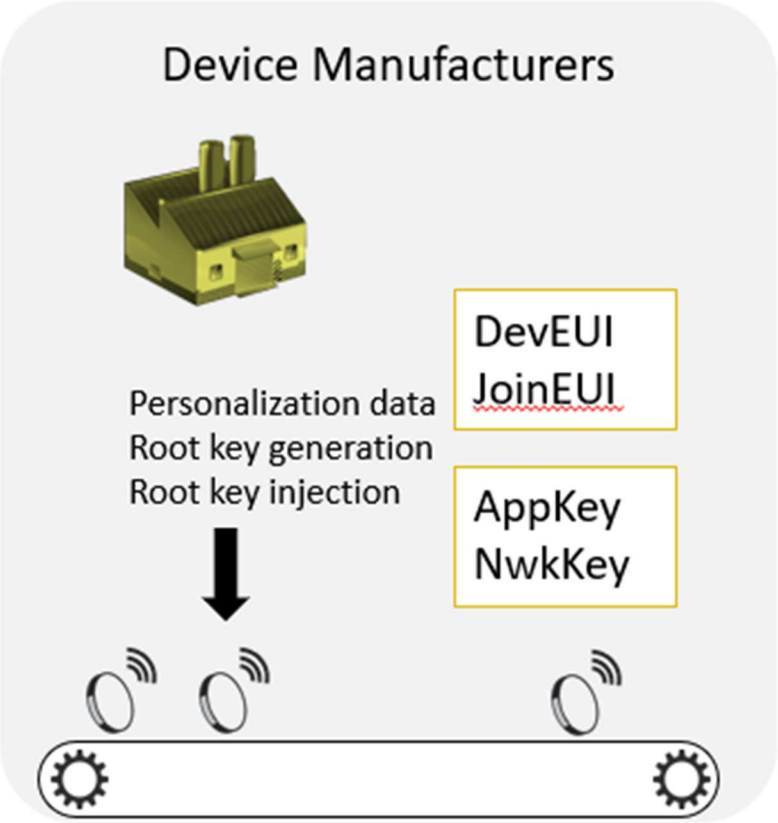

Figure 1 - LoRaWAN end-device personalization

DevEUI is a global end device ID in the IEEE EUI64 address space that uniquely identifies the end device. It is used to identify a device roaming across LoRaWAN networks and is selected by the manufacturer from the range assigned by IEEE.

JoinEUI is a global ID in the IEEE EUI64 address space that uniquely identifies the join server that is able to assist in the processing of the join procedure and the session keys derivation. It is therefore selected depending on which join server stores the device root keys, in the range assigned by IEEE to the join server supplier.

The root keys **NwkKey** and **AppKey** are AES-128 keys specific to the end device. They should be generated or derived randomly with high entropy so that they cannot be guessed by observing other devices of the same batch.

The LoRaWAN root keys are the root of security of the entire system. One critical element of the device manufacturing floor is thus the key injection step. In most secure manufacturing lines, keys are generated in Hardware Secure Modules (HSMs), that are tamper-proof physical appliances protecting secrets from physical and remote access. These appliances are connected securely to the appliance injecting the root keys into the device, so that keys are not exposed, even during manufacturing. It is also possible to design secure remote protocols to

generate and inject cryptographic secrets into the device via other connectivity layers than LoRaWAN.

For applications requiring a high level of security comparable to SIM cards, devices can be equipped with secure elements. They are tamper-proof subsystems soldered on the device holding root keys injected by secure element manufacturers under high security and deriving session keys and MIC separately from any application code running on the main processor. They ensure no spying or hacking is possible when a device is physically accessible to attackers.

The final step of personalization is the selection of a join server that has activation agreements with the target network server operator(s). Each join server can expose one or multiple JoinEUI that uniquely identify it across all LoRaWAN networks. This value must be put in the device during the personalization phase so that the join server can be identified during activation. This join server must be known by all target network server operators, so they can forward the join request to it.

How are LoRaWAN End Devices Provisioned?

Before activation, devices must be provisioned out-of-band into the three LoRaWAN network elements: join server, network server, and application server.

Once the join server supplier is identified, keys are shared with it directly or indirectly (through Key Management Systems, or KMS). This step must be secured at the transport level and at the storage level on both sides. The join server (and KMS) also often make use of HSMs to secure the storage and even the transport of these secrets.

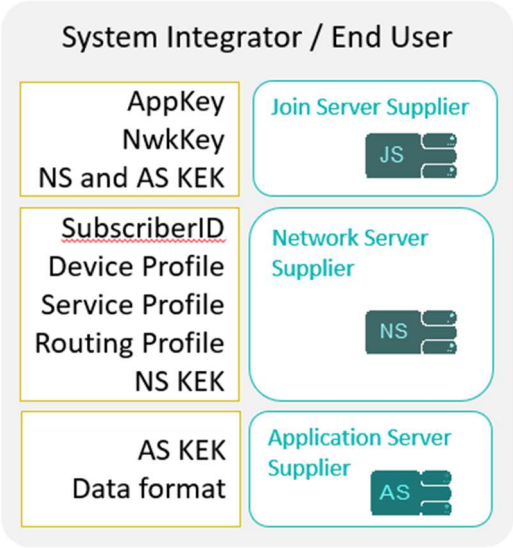

Figure 2 - Provisioning information in LoRaWAN network element

Join servers are also provisioned with Key Encryption Keys (KEK) that are used to transport the session keys derived during device activation to the network server and the application server. These keys are AES-128 keys and are shared between the join server and network/application server securely. Sharing takes place in a similar manner to that of the device root keys previously described, but depend only on network elements, not devices.

The join server must also know the home network server for the home discovery procedure, so it is provisioned with its NetId for roaming use cases.

Devices must also be provisioned on a home network server. This server holds the primary connectivity for the device and can interact with visited network servers in case of roaming. The home network server must be provisioned with all necessary information to manage connectivity: device profile with radio boot parameters and supported MAC commands, connectivity plan, routing profiles to application server(s).

Finally, devices are provisioned in an application server. This server runs the end-user application(s). It should be provisioned with transport- and application-level information: KEK is used to transport the AppSKey, device data decoder, and any other business logic. The LoRaWAN standard specifies how binary data and the AppSKey are delivered to the application server. However, the rest of the security is out of the specification scope.

How are LoRaWAN End Devices Activated?

In the final step, device activation, the end device is switched on and activated. Its first action is to perform a join request procedure to assign a temporary security session. The join procedure has been designed for low-power, low-throughput efficiency and high security.
A minimum set of parameters is transmitted via a join request uplink message: DevEUI to identify the device uniquely, joinEUI to identify the join server uniquely and a DevNonce counter for security.

Size (bytes)	8	8	2
Join-request	JoinEUI	DevEUI	DevNonce

The network answers with a set of radio parameters, a DevAddr to identify the device uniquely on the radio containing the home network NwkId (a subfield of the 24-bit NetId) and a JoinNonce counter for security.

Size (bytes)	3	3	4	1	1	(16) Optional
Join-accept	JoinNonce	Home_NetID	DevAddr	DLSettings	RxDelay	CFList

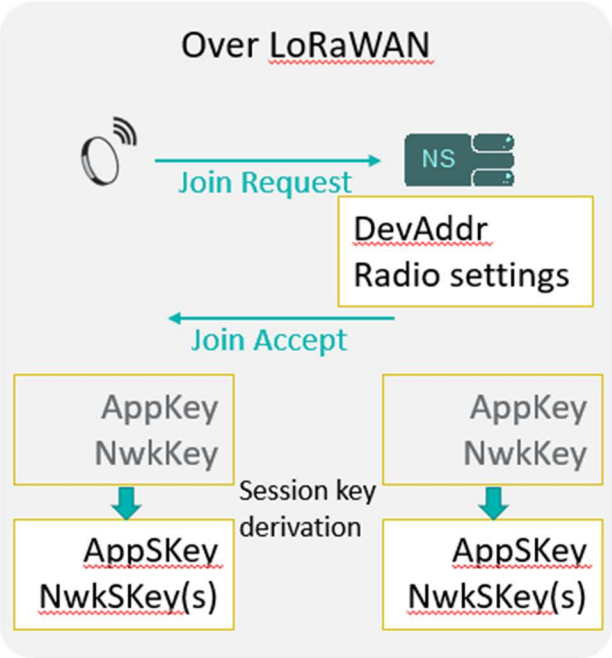

Figure 3 - LoRaWAN End-device activation: Join procedure radio messaging

The purpose of the DevNonce and JoinNonce counters is to protect the procedure against replay attacks. These values ensure that the messages are never identical for two different procedures, so that we cannot record a message and replay it later without the end device or network refusing it. Defining the nonces as counters allows the device (resp. network) to store the last JoinNonce (resp. DevNonce) used during the previous procedure and easily verify that these values haven't been re-used: The join server simply verifies that the new value is strictly greater than the previous one and has been incremented by one.

Again, LoRaWAN is using very simple mechanisms such as monotonically increasing counters to simplify implementation and allows very low-cost devices to comply with the standard.
Full end-to-end security is achieved through session key derivation that happens in parallel on both sides: end device and join server. The AES-128 device root keys were pre-shared securely between the end device and join server, so the root key derivation is secure as long as its execution and storage of resulting session keys are protected on both sides. Again, this is implementation-specific and most secure implementation relies on secure elements on the end device side and HSM on the network side.

LoRaWAN also allows device activation away from home, that is to say on a visited network. In this case, scalability is achieved through the JoinEUI identifier, which allows dynamic discovery of the home network.

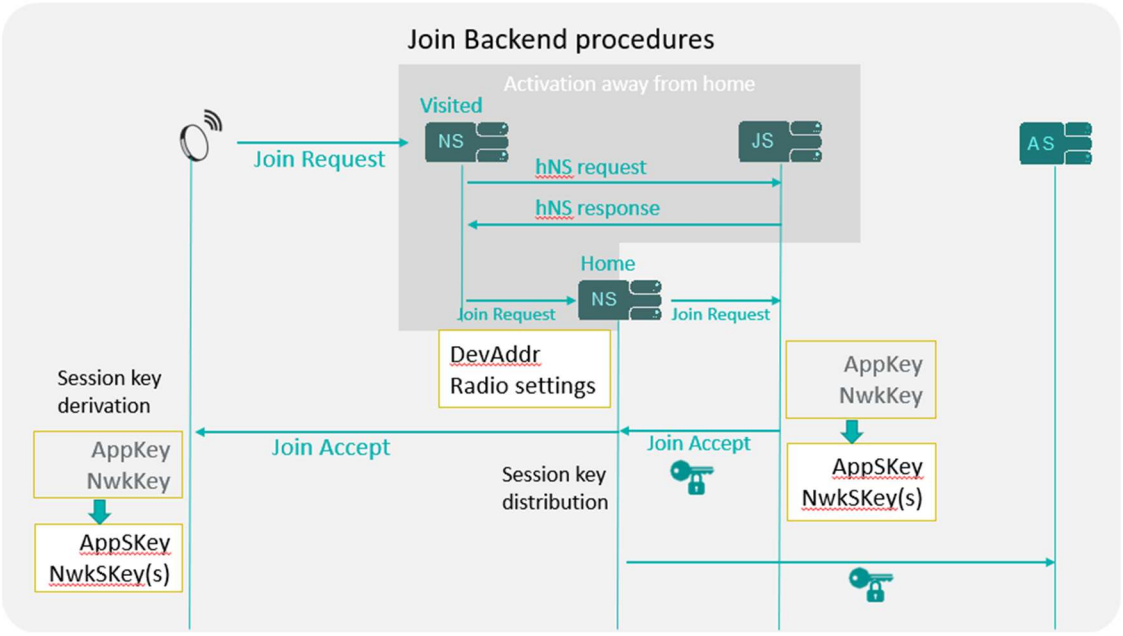

Figure 4 - LoRaWAN end-device activation: Join procedure backend messaging

A visited network needs to have home discovery agreements with the home join server that provides JoinEUI mapping to the join server address and credentials for secure transport layer negotiation with the join server. This is further facilitated by the use of the LoRaWAN roaming hub on LoRaWAN Roaming. In this case, the visited network server discovers the home network server through the Hide and Seek (HNS) discovery procedure. For each device, the join server records the provisioned home network server during provisioning and sends it back to the visited networks. The join request is then simply forwarded to the home network server, and the rest of the procedure is similar to the OTA Activation at Home procedure.

At the end of the activation process, the temporary session keys are distributed by the join server to the other network elements: NwkSKey(s) to the home network server (and forwarding network server for LoRaWAN 1.1 devices) and AppSKey to the application server. This is secured by the KEK previously exchanged during provisioning and standardized in the LoRaWAN backend interface.
Finally, the home network server ensures data integrity and authenticity using the NwkSKey(s). The application server receives data and the AppSKey encrypted with KEK together. The network server, and the join server itself when an HSM appliance is used, never have access to an unencrypted AppSKey, which guarantees end-to-end confidentiality from end device to application server.
Similarly, on the device side, session keys can be derived from root keys in a secure element or secure OS environment and used within these secure environments to guarantee data integrity, authenticity, and confidentiality at the device level as well, making it resistant to physical attacks.

Secure, Scalable and Power Efficient LoRaWAN Activation

The use of symmetric cryptography and prior key sharing between the device and the join server enables an extremely power efficient and network efficient activation procedure. In traditional networks, we use asymmetric cryptography to avoid the need for prior key sharing, but this method pushes the complexity back on the radio, as the device and join server equivalents must negotiate a secure session before being able to communicate. Application-level security relying on public key signatures can span multiple frames, consume network capacity, and drain the end device battery.
Similarly, network servers (or LoRaWAN roaming hubs) are provisioned with join server connectivity settings based on JoinEUI values. Such provisioning simplifies procedures and allows flexible architectures where many different join server suppliers can be connected to network servers, such as private enterprise join servers.

LoRaWAN activation covers end-to-end security, from parallel session key derivation in the device and join server to key distribution over multiple roaming networks, securing integrity, authenticity, and end-to-end confidentiality of the device data up to the application server.

About the author
Raphael Apfeldorfer is currently serving as Advanced Technology Director at Actility and is leading the introduction of new ThingPark IoT products to the market. He has over 19 years of experience in the wireless communication industry ranging from cellular networks and devices to IoT networks. Prior to his current position, Raphael has been leading NVIDIA European Application Engineering team, helping numerous customer products based on NVIDIA/Icera cellular modems reach the market in phones, tablets, automotive telematics, 3G/4G datacards and USB dongles. Prior to NVIDIA/Icera, Raphael started his career in 1999 at Nortel Networks where he led embedded signal processing development of first 3G radio access networks.

About the company
Actility connects the industrial internet of things. Our IoT connectivity platform, tools, and fast-growing ecosystem enable our customers to create IoT solutions that transform business, industries and processes. The ThingPark LPWA platform connects sensors gathering data to cloud applications on any scale, from global or national networks to secure on-campus enterprise solutions, managing devices, data flows and monetization. Our value-added applications and business services enable roaming, device software update, geolocation and smart grid. Actility is at the heart of a thriving customer ecosystem, connecting solutions partners, supporting developers and device makers preparing their LPWA product for market, and providing an e-commerce Marketplace offering global distribution to solution providers. Actility co-founded the LoRa Alliance and continues to pioneer LPWA networking technology.

https://www.actility.com/

4.4 Building Secure IoT Applications: From Design to Implementation

By Jan Stegenga, Onethinx b.v., Zwolle, The Netherlands

Introduction

New applications of LoRaWAN rely on its promise of low cost, large-scale sensor data to make decisions that are critical to a company's operations. More important, these decisions may be partly or fully automated, so data integrity and with that, security, are of utmost importance. The latter entails that no party other than the owner can access the root keys, device software, or sensor data.

Security is a must-have for new IoT applications to become successful. The LoRaWAN specification is set up for secure transport, with content encrypted by application session keys and network session keys (LoRaWAN 1.1). While such transport tackles general security for servers and gateways, two important issues remain, which the LoRaWAN specification does not include. First, there is the secure management of root keys. These need to be accessible to the join server and stored on each individual device. Any party that knows these keys can access decrypted data and build devices that mimic genuine devices. Second, without secure operation of the physical device, it's possible to read the software and settings (such as root keys), alter the device's functionality, and insert spurious data. We address these two issues in this paper. First, we describe a process that will guarantee the secure storage and distribution of these keys during manufacturing as well as deployment. Second, we describe how to create a secure device at the hardware and firmware levels, where the recent release of the Arm® Platform Security Architecture (PSA) has spurred secure functionality on SoMs, SoCs, and microprocessors.

Secure Provisioning and Manufacturing

The LoRaWAN 1.02 specification defines three 128-bit security keys. The application key (AppKey) is known by the device and by the application. When a device joins a network, an application session key and a network session key are generated. The network session key is shared with the network for authenticating the device, while the application session key that is used for encrypting the data is kept private. These session keys will be used for the duration of the session. As such, in LoRaWAN 1.02, the root key is the AppKey. In LoRaWAN 1.1, there are two root keys: AppKey and NwkKey. All root keys must be stored securely in the LoRaWAN end device and on the network side (join server) as well because the encryption is symmetric. See [TTN, Jialuo Han and Jidong Wang] for more information.

In the following model, the Key Management System (KMS) will issue the root keys at the time of manufacturing and make them available to the join server so that session keys can be derived from them when a device attempts to join a network. Apart from the KMS, no other party has access to root keys: the LoRaWAN service provider only has access to network session

keys for device authentication, and the device owner only has access to the application session key to retrieve sensor data. The architecture of a LoRaWAN network with KMS and factory-key provisioning is shown in Figure 1.

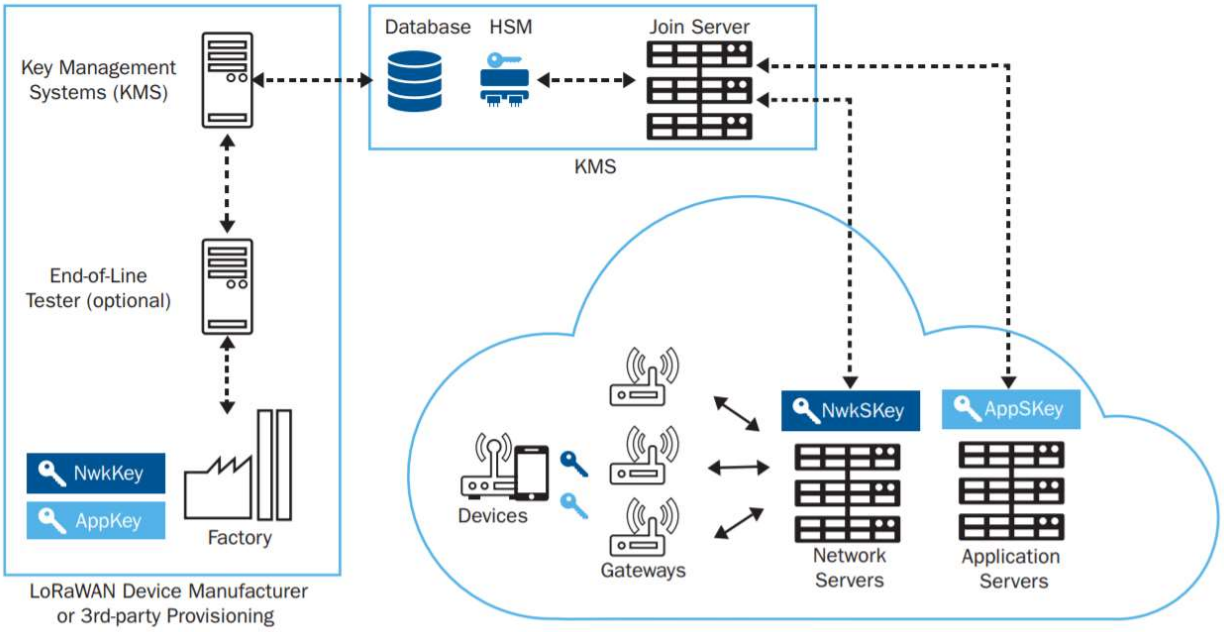

Figure 1: Key management scheme. A KMS is responsible for issuing root keys (provisioning) and providing them to a join server in the join procedure.

Device key provisioning can be done either during device manufacturing or in the field, as devices are registered to join a LoRaWAN network. Only after (two-factor) authentication of the party doing the provisioning will the KMS supply the root keys. If applicable, the devices will immediately be registered to a device owner and LoRaWAN service provider.

During a join request, network service providers access the KMS/Join Server using a specification of the LoRa Alliance. Again, such connections use end-point authentication, integrity, replay protection, and confidentiality. With key management separated from the owners and service providers, the KMS can be optimized for security and availability. This way, roles are distributed among parties and each party only has access to the information it requires to carry out its role.

Secure Nodes

The developer of the end node device must ensure that root keys can be stored on the device. To streamline development of secure software and hardware, Arm has introduced its Platform Security Architecture (PSA), as noted above. The PSoC®6 family of microcontrollers (MCUs) from Cypress are among the first to support PSA. The PSoC MCUs provide three levels of hardware-based resource isolation, the highest level of isolation defined by PSA. The two-core System-on-Chip offers the possibility of a secure core, physically separated from a user application running on the other core. In addition, PSoC6 offers secure element functionality that can be used to build a chain of verified and secured applications. This chain of trust (CoT) protects the application(s) from being altered by calculating and storing hashes over code blocks and signing these with private/public key pairs.

Onethinx has worked on the implementation of a PSA-level secure system that is extendible by third party developers. This serves three objectives:

1) A locked-down LoRaWAN stack on a module with an integrated antenna makes the certification process much simpler.
2) The Onethinx code is copy-protected.
3) Third-party developers can build their application on the SoC and secure it.

The PSA approach defines that an MCU must start up in a trusted manner. This is ensured by a section of Read-only Memory (ROM) code that will calculate a hash over the next section of flash executable code and compare that hash with one that is stored in a write-once memory location (a hardware register called eFUSE). The processor will start executing code from the next section only if the hashes match. This is the trusted root that is built upon by subsequent applications.

A chain of trust is formed by one or more applications that execute sequentially, where each validates the next application in the same way as the root element. The validation is done by ROM functions, the stored hash is signed using a private key of the code developer, and its memory is protected from alteration in the sense that if it is altered, the code will not be executed. There are several methods to protect a section of flash memory from being read or written to as well. These are used to protect the memory location of the root keys from being accessed by unauthorized applications/resources.
If the chain of trust is to be extendible by third-party developers, each application should be signed with the respective code developer’s private key. An application programming interface (API) for the secure storage of the corresponding public key must be made available to functions just like the eFUSE in the secure root.

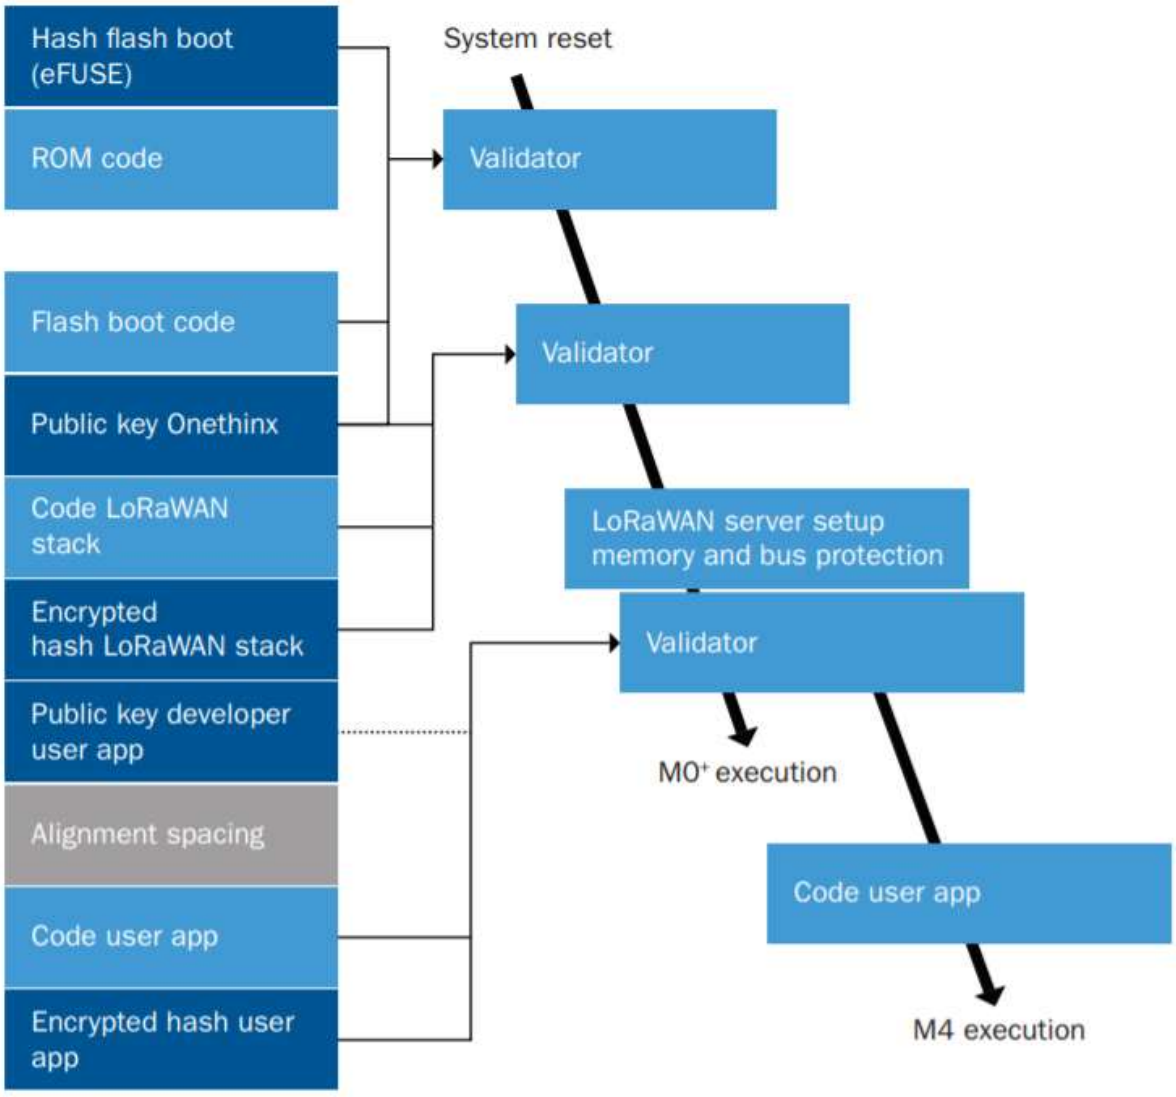

Figure 2: Chain of Trust from reset to user application start. Each component builds upon the chain by validating the code of the next app before executing it.

In Figure 2, the sequence of links is illustrated. The left side shows the sectioning of flash, the right side shows the execution order and what information is used for validation. The PSoC6 has two Arm cores, an M0+ and an M4, which are connected by the Inter Processor Bus (IPC) that also connects other components (flash, IO, etc.). The first link is the flash boot. It runs on the Arm Cortex® M0+ core, which is designated as the secure core and is validated by ROM using the hash stored in eFUSE. The second link is the Onethinx LoRaWAN stack. It provides LoRaWAN functionality in a server-client style over the IPC, as the PSA recommends. Next, it sets up memory, peripheral, and bus access protection that isolates the M0+ (the secure core) from the M4 (regarded as insecure). The CoT is broken when a calculated hash does not match a (decrypted) stored hash. This indicates that the flash code has been altered in some way, which is a potential security breach, and will cause the chip to enter a dead state.

The user application runs on the M4 core and uses IPC calls to initialize the radio link, set and store LoRaWAN join parameters, and send and receive LoRaWAN messages. The true number generator is used for generating the random number in LoRaWAN join messages required to prevent replay attacks. Physical separation of the user app and the stack ensures that mistakes in the user code will never interfere with the proper functioning of LoRaWAN communication, thus speeding LoRaWAN certification. In addition, the LoRaWAN join keys can also be stored safely.

Reading, modifying, and executing code on the M0+ must be prevented. Following the setup listed in AN22111:

- Sections in Random-access Memory (RAM) and flash are defined and protected for usage by either M0+ or M4
- Access to I/O that is used for LoRaWAN is restricted
- Access to potentially harmful IPC calls are restricted
- The debug port on the M0+ (M0+-DAP) is disabled, and the programming port is not allowed to access the regions associated with LoRaWAN/M0+

Since all resources are available through the IPC bus and mapped to memory locations, controlling access to the IPC and memory locations based on the requestor's attributes is a very powerful mechanism. The M0+ and M4 cores are IPC bus masters and can have attributes such as protection context, secure designation, or elevated rights. In the LoRaWAN stack, the bus masters are given different protection contexts, which are used by the protection units to control access. Hence, a call from the M4 (PC=2) to access the peripherals of the radio chip (the peripheral protection unit allows only PC<2) will be disallowed. The same mechanism is used to separate RAM and flash. The linker script ensures that code and variables are put at the pre-defined locations.

Several functions, such as flash writes, are implemented through IPC function calls and are always carried out on the M0+. Restriction of these functions is done in software by intercepting the callback function handling IPC calls and providing access depending on caller, function, and memory location.
The debug ports can be disabled in the access registers of the PSoC6. There are three access registers, two for the life cycles 'normal' and 'secure' and one for when the MCU enters 'dead' mode. The LoRaWAN stack writes minimal restrictions to all registers to ensure that the M0+ debug port remains disabled and the programming ports cannot overwrite critical sections. The chip is set in the 'secure' life cycle stage.

Continuation of code development on a system that has been partially locked-down is quite unique to the LoRaWAN stack. An unknown developer must be able to program and debug his/her application and should have as many of the PSoC6 resources available as is securely possible. To secure their code developers can write their public key to a secure part in memory by calling a function provided by the LoRaWAN stack. This can be done once, and from that

moment onwards the user app code must be signed with a hash that is encoded by the developers' private key to be executed.

Summary

It is both possible and necessary to implement security in the entire LoRaWAN chain and for the IoT overall. Significant improvements in the security of LoRaWAN have been implemented in the LoRaWAN 1.1 specification. Key management can be separated from owner and service provider such that each party only has the keys for the part required to fulfill its role. Microcontrollers with secure elements for the secure storage of keys and securing the whole software are available. With the Arm Platform Security Architecture, an ecosystem for streamlined implementation of security functionality in both hardware and software is being built. Real-time Operating System (RTOS) support for these is under development and is expected to provide a mechanism for secure firmware updates as well.

About the Author

Jan Stegenga (1979), MSc, PhD, has a background in electrical engineering. He specializes in adaptive, data-driven, modelling of time series data. He translates such models to real-time implementations in resource limited hardware such as microprocessors and FPGAs. For Onethinx he develops the security aspects of the Onethinx LoRaWAN module and implements customer specific solutions.

About the Company

Onethinx b.v. (Zwolle, The Netherlands) specializes in full service IoT development solutions, with a focus on LoRaWAN networks. They follow pragmatic procedures resulting in attractive and user-friendly applications. They offer off-the-shelf and tailor-made applications based on their overall conceptual and technical approach. Their programmable LoRaWAN module with secure stack and integrated antenna is their flagship product. It goes beyond the market standard in terms of ease of implementation and versatility.

https://www.onethinx.com/

LoRaWAN ARCHITECTURE

5.1 LoRaWAN ARCHITECTURE OVERVIEW

By Semtech

A LoRaWAN network is always implemented in a star topology, as shown in Figure 1. Unlike a mesh topology, a star topology is ideal for power-constrained (i.e., battery-operated) end-nodes (sensors) because each end-node transmits only its own messages, so no battery life is consumed relaying messages from surrounding end-nodes. A LoRaWAN network consists of one or more LoRaWAN gateways that are all connected to one central network coordinator, a network server (NS).

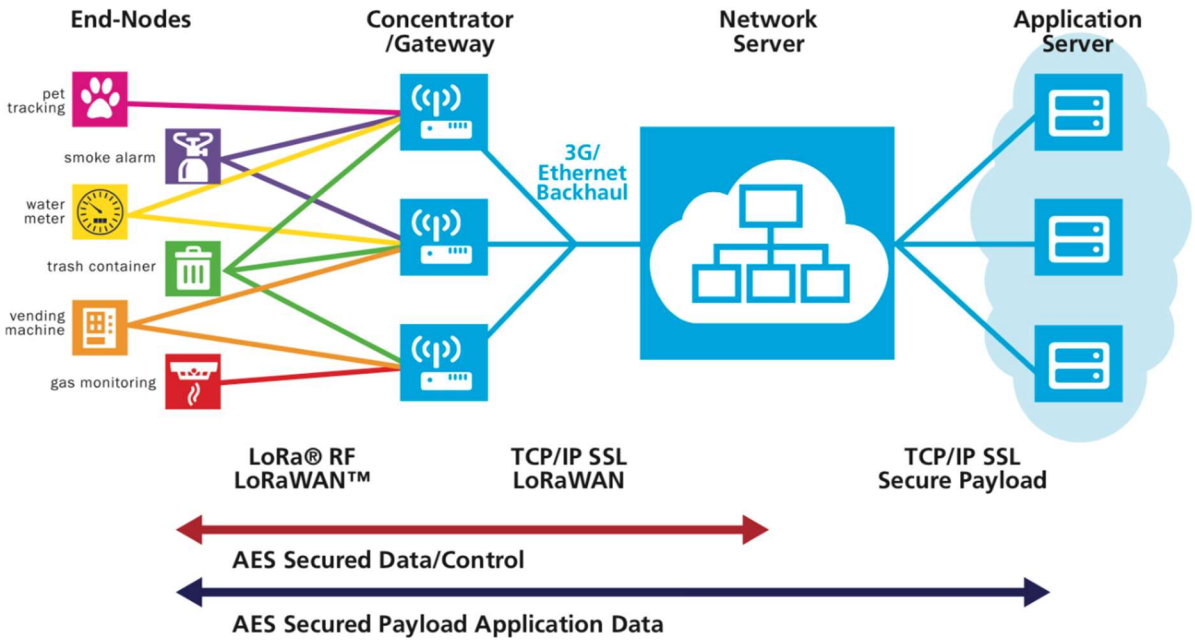

Figure 1: LoRaWAN Star Network Topology
© Semtech Corporation. Reprinted with permission

Unlike cellular base stations which have a high level of hardware and software complexity, and therefore a high cost, LoRaWAN gateways are basic protocol bridges with a much lower cost point. Each gateway receives LoRa modulated radio messages from all LoRaWAN end-nodes within radio distance. Every received LoRaWAN frame with a correct cyclic redundancy check (CRC) code will be forwarded to the NS, encapsulated in an IP frame. Gateways can be connected to the NS over Wi-Fi, hard-wired Ethernet, or even a cellular connection. The gateway is, in essence, a bridge between LoRaWAN and IP. A managed LoRaWAN network typically consists of both macro (64 channels) and picocell (8 channels) gateways. The macro cell solutions provide citywide (broad) coverage, while the picocell gateways allow for increased network capacity in dense areas. An increase in the number of gateways in a specific area will

typically prolong LoRaWAN end-node battery life in that same area because the distance between end-nodes and gateways is shorter. This will in turn allow LoRa radio packets to be transmitted with lower spreading factors, which require shorter Time on Air (TOA). For example, a LoRaWAN frame with an 11-byte payload transmitted in a LoRa radio packet with SF7 has a TOA of 61ms, whereas that same payload would require 371ms when transmitted at SF10 (as can be seen from Table 1). Another benefit of densifying a LoRaWAN network with picocell gateways is the ability to assure coverage in hard to reach areas such as basements.

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

Table 1: LoRaWAN Spreading Factors with TOA for an 11-byte Payload

In the network architecture pictured in Figure 1, the NS actively manages the LoRaWAN network. It coordinates incoming messages (LoRaWAN frames) from all the different end-nodes as well as network commands (MAC frames) between all the gateways and the application servers. Below is a list of the major management functions implemented in the NS:

- **IoT device Over-the-Air Activation (OTAA).** The NS manages all requests from end-nodes to access the network (Join Requests), informing each end-node which set of 8, 16 or all 64 channels to operate on, what spreading factor (SF) to use, and what power level to transmit.

- **Data de-duplication.** The NS deletes duplicate radio messages received from the same IoT devices from different gateways.

- **Dynamic data link (DL) LoRaWAN frame routing.** Selects the gateway best suited to connect with each LoRaWAN end-node to receive DL messages.

- **Adaptive Rate Control (ADR).** ADR's main goal is saving LoRaWAN end-nodes' battery power. By having the end-nodes closest to a gateway transmit using the lowest SF, their time on air is minimized, thus prolonging their battery life. More distant sensors transmit at a higher SF. A tradeoff is made between battery power and distance, as a higher SF allows for a gateway to connect to a sensor further away.
- **Network congestion.** The NS can instruct individual end-nodes to connect to a different gateway by changing their channel plan(s).

- **Forwards all application data to the right application server.** When a LoRaWAN end-node requests to join the network using a join request MAC command, the end-node transmits a specific application code (AppEUI), a unique code indicating which application server its data needs to be forwarded to. Once the end-node receives a join accept MAC command, all further frames containing user (application) data will be delivered to its designated application server.

- **Administration.** General network administration, network provisioning, and reporting functions.

The NS has control over key operational and network settings in each LoRaWAN end-node: Transmit Power Level, SF, Channel Plan (i.e., which set or sets of 8 channels out of the 64 available 125KHz uplink channels), receive window timing, etc. Control over the above-mentioned parameters provides the NS the ability to optimize both network performance and end-node battery life. Critics might argue that one of the downsides of a star topology is that when the NS goes down, no data traffic is possible. A real-world managed LoRaWAN network implementation would of course have a redundant NS implementation to cover such a scenario.

LoRaWAN has been deployed by 120 public operators in more than 140 countries to date, and the technology is enabling products that are smarter, cleaner, and more effective for managing resources. For end users, LoRaWAN delivers hard, real-time metrics that can be used to shape greater operational advantages. For designers, the decision to use simple, long-range, and low-power LoRaWAN is driven by the ambition to create systems that are innovative and strongly efficient.

https://www.semtech.com/

LoRaWAN DESIGN SERIES

6.1 LoRaWAN Roaming

By Alper Yegin, Director of Standards and Advanced Technology Development at Actility, Vice Chairman of LoRa Alliance

Roaming is the ability to use radio coverage provided by foreign networks when a device is outside the coverage of its home network. This concept is well-established and commonly used with wireless access technologies, including Wi-Fi and cellular ones. Roaming is also being implemented by LoRaWAN networks. LoRaWAN roaming follows the established concept at a high level but differs from its legacy counterparts in some aspects, as the IoT demands cost-effectiveness and high scalability.

There are three basic cases in which an IoT device may find itself roaming into another network:

- Mobile devices crossing network borders, such as trackers.

- Fixed devices of a multinational enterprise. It is likely that the single public operator contracted by the enterprise (i.e., home operator) does not have radio coverage in all countries where the enterprise has businesses. Instead, the home operator relies on a network of international roaming partners to provide the required coverage for its enterprise customer devices. Using operators with partial network coverage within the same country also falls under this case.

- Any kind of device (fixed or mobile) taking full advantage of radio coverage provided by multiple networks even when it is under the coverage of its home network.

The last type of roaming is a very unique feature provided by LoRaWAN. An uplink frame sent by a device can be demodulated by any LoRaWAN gateway within the radio coverage, irrespective of which network it belongs to, and sent to the device's home network. In this way, multiple networks can collaborate together to provide extended and densified network coverage. This feature benefits the home network, which otherwise may not be able to reach the device using its own radio network. The device benefits from this collaboration, as the additional networks cause gateway densification, which translates to reduced transmission power and repetition required for successful transmission, and therefore longer battery life by using the Adaptive Data Rate (ADR) feature of LoRaWAN[1]. Finally, the roaming partners of the home network (i.e., foreign networks) enjoy a reduction in interference, which would have been caused by the device increasing its transmission power and spreading factor when its uplink frames were not directly received by its home network.

[1] LoRaWAN Specification, Version 1.1, LoRa Alliance, October 2017.

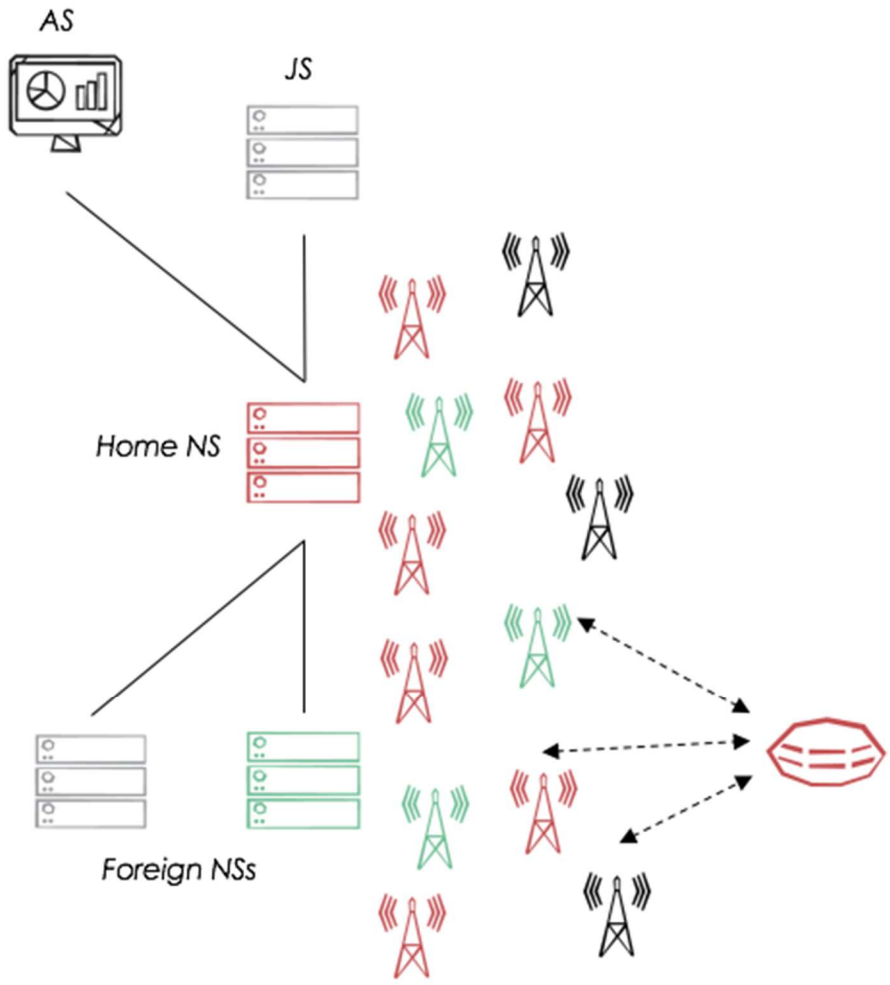

Figure 1. LoRaWAN Passive Roaming.

The type of roaming based on collaborative reception is called Passive Roaming (see Figure 1). The passive nature of this type of roaming comes from the fact that roaming is totally transparent to the device. As such it does not require any special handling by the device and it can be used with devices implementing any version of the LoRaWAN specification.

At the core of Passive Roaming are two essential attributes of LoRaWAN: First, gateways are stateless, and, second, the ecosystem is open for multiple networks to coexist. Cellular IoT technologies like NB-IoT are built on top of architectures with stateful base stations, an approach which prevents the macro-diversity needed for this level of collaboration. While this technical obstacle does not exist for SigFox, another LPWAN technology, its locked ecosystem does not leave any room for multiple networks in the same geography. Therefore LoRaWAN single-handedly enjoys this feature that further accelerates infrastructure cost reduction as required by IoT.

The other flavor of roaming supported by LoRaWAN is Handover Roaming. In Passive Roaming, the foreign networks whose gateways are used for extending and densifying radio coverage act as mere radio network extensions to the home network. The LoRaWAN link-layer of the device stays connected to its home network. The MAC commands sent by the network are generated by the home network, and the ones sent by the device are consumed by the home network. In the case of Handover Roaming, however, the link-layer control is handed over from the home network to a foreign network.

Handover Roaming requires devices to perform specific signaling which is defined in the LoRaWAN 1.1 specification. Unlike its counterpart, Handover Roaming cannot be used with devices implementing LoRaWAN specifications earlier than version 1.1. On the other hand, for optimum network utilization, the home network better knows the channel plan of the foreign networks, so it can adjust the channels used by the device to maximize uplink frame reception by a given network(s) when using Passive Roaming. Both Passive and Handover Roaming rely on an uplink frame to be received by both home and foreign networks. For that reason, some level of channel overlap among roaming partners is required.
Implementation of roaming requires networks to implement the LoRaWAN backend interfaces specification[1]. This specification describes the signaling among the network servers and join servers that are operated by separate organizations. Roaming requires a network server that receives a frame to deliver that frame to another network server, and interoperability of this delivery is achieved by the aforementioned specification.

Among the two flavors of roaming, Passive Roaming is the one that is implemented and deployed first. The power of collaborative reception, the absence of dependency on any specific version of LoRaWAN implementation on the device, and sufficiency to get roaming up and running are the reasons we are seeing the industry start with this type of roaming. Nevertheless, both Passive Roaming and Handover Roaming can be used in conjunction (see Figure 2). A device can be subjected to both types of roaming at the same time. For example, a device that is far away from its home network may be under the coverage of several networks with which the home network may not have a roaming agreement. In that case, the power of Passive Roaming cannot be fully exercised. If the home network happens to have a roaming agreement with one of those networks, it can set up Handover Roaming with that network, which in return can set up Passive Roaming with the others in its vicinity. Handover Roaming is used to extend the reach of Passive Roaming in this scenario.

[1] LoRaWAN Backend Interfaces Specification, Version 1.0, LoRa Alliance, October 2017.

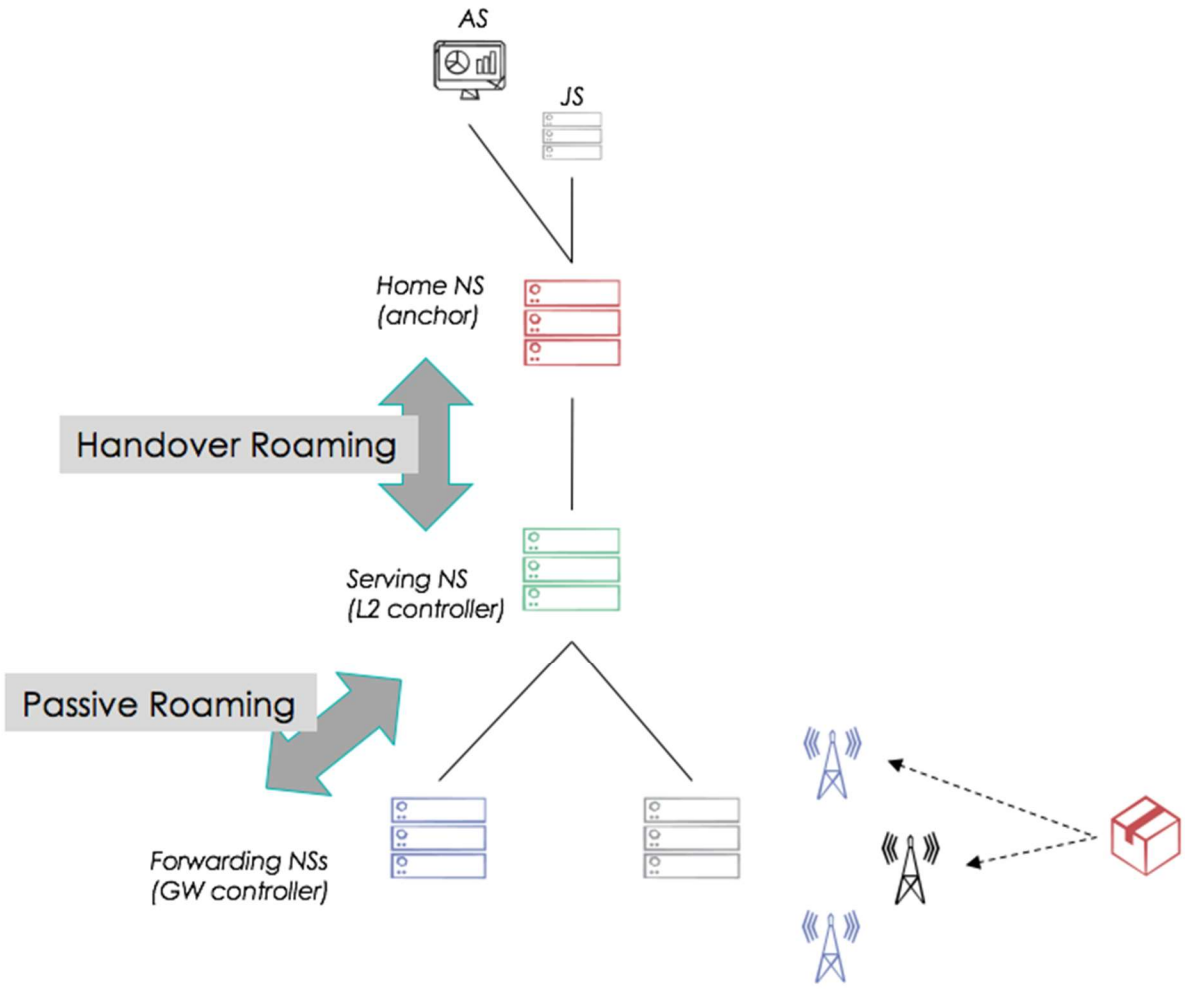

Figure 2. Use of Passive Roaming and Handover Roaming together.

Another essential difference between LoRaWAN and other LPWAN technologies in the area of roaming is its applicability: LoRaWAN roaming can be used between public networks, between public and private networks, and even between private networks. For example, two neighboring cities with their own private networks (i.e., not selling connectivity to 3rd parties) can decide to allow each other's devices to roam in. LoRaWAN with its open and collaborative nature presents itself as a rich toolbox ready to be used in any way demanded by the market.

Roaming between two networks requires setting up a roaming agreement and interconnecting the networks with each other. As more and more networks decide to enable roaming with each other, this leads to the number of agreements and interconnections growing beyond a manageable point. This is where the well-established roaming hub concept (see Figure 3) comes into the picture. A mesh-interconnect among several networks can be turned into a star-interconnect via a hub when each network connects to the hub. We are already seeing deployment and utilization of LoRaWAN roaming hubs. What differentiates LoRaWAN roaming hubs from the ones used for the other technologies (e.g., Wi-Fi, cellular) is the low-cost and high scalability of the former.

A LoRaWAN roaming hub implementing Passive Roaming presents itself as a transparent network server proxy as it is inserted between two network servers. Therefore implementation of the roaming hub can be achieved by following the interfaces defined in the LoRaWAN backend interfaces specification.

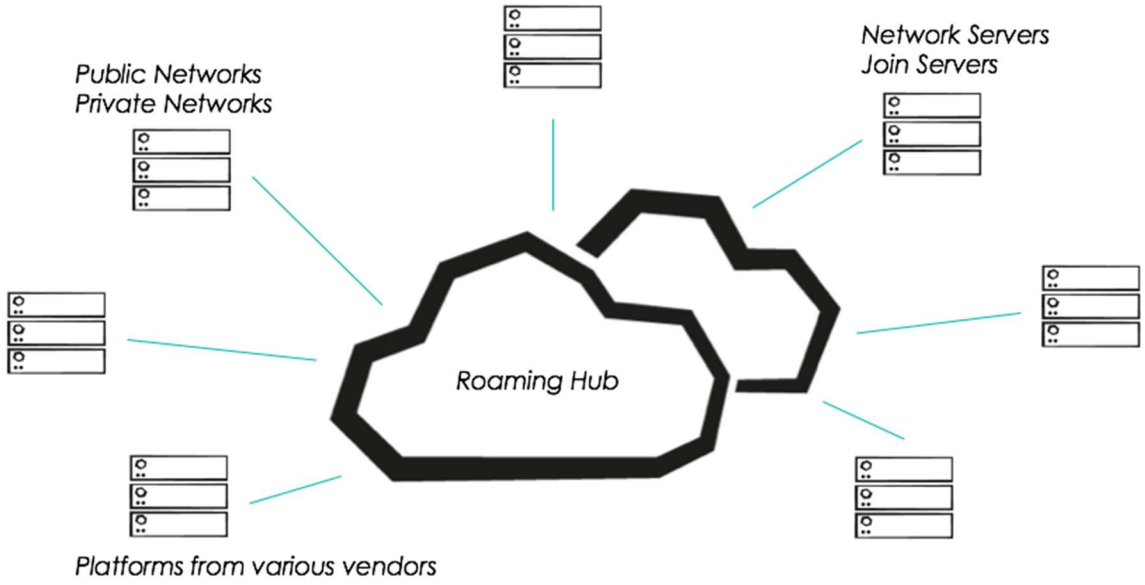

Figure 3. LoRaWAN Roaming Hub.

Roaming is a directional service. A network may decide to allow devices from another network to roam into its own radio coverage, while the other network may or may not reciprocate this service. This is just a mutual roaming policy matter. Each network keeps track of the amount of roaming service provided to the other, and the basic metric being used is the frame counter. The number of uplink and downlink frames delivered for a roaming partner is included in the generated Usage Data Records (UDR). UDRs are needed at least for network diagnostics and also potentially for supporting the commercial agreements between the roaming partners. Metrics other than frame counters may be defined in the future in order to address evolving business models. LoRaWAN networks and their devices are identified based on the so-called NetIDs. These are 24-bit identifiers assigned by the LoRa Alliance to its members. A network server assigns each of its devices an address (DevAddr) using an identifier derived from the NetID. That allows any foreign network that receives an uplink from a roaming device to identify the home network of the device and forward the frame towards that network. One or more foreign networks receiving an uplink frame and able to forward the frame to its home network without having any context about the device is at the core of the Passive Roaming feature. Therefore, it is very important that the device addresses are based on the NetIDs assigned to the network by the LoRa Alliance. Assigning arbitrary device addresses is highly discouraged as it prevents the device from ever performing passive roaming in its lifetime. Similarly, the join server identifier setting on the device (AppEUI/JoinEUI) needs to point to a server that can

authenticate the device. Using an arbitrary identifier instead prevents the device from joining the network when it is under the coverage of a foreign network. Using DNS with LoRaWAN-specific extensions for facilitating IP address resolution of network servers and join servers is also one of the tools provided by the LoRaWAN specifications.

LoRaWAN specifications and implementations provide the essential enablers for widespread roaming. At the time of writing this article, we are already seeing networks setting up roaming with each other and getting ready to benefit from the unique opportunity provided by LoRaWAN. LoRaWAN shines among competing LPWAN technologies thanks to its uniquely open and collaborative nature, which translates to power-efficient connectivity for the devices, and low-cost and scalable networks as demanded by IoT use cases. In order to reach that potential, we would like to urge network operators, big and small, private and public, to seek partners to set up roaming. Roaming should be used by default. Network owners should not look for reasons to roam, but instead, they should look for reasons not to roam with the others. Roaming acts as the glue uniting networks towards providing global coverage and creating a win-win situation while still respecting the commercial priorities of the players involved.

About the Author

Alper Yegin is a technology architect involved in research, design, and standardization of IoT and mobile technologies. He is currently serving as the Director of Standards and Advanced Technology Development at Actility, co-chair of the Technical Committee and vice-chairman of the LoRa Alliance. Prior to his current post, Alper has worked for Samsung Electronics Research Center where he led the design of 5G IP mobility, 4G WiMAX security, and ETSI M2M security. He made significant contributions to the design and standardization of networking technologies including Mobile IP, IPv6, Zigbee IP, and PANA during his tenure at Samsung, DoCoMo USA Labs, and Sun Microsystems. He has been actively involved in international standards organizations such as LoRa Alliance, IETF, ETSI, 3GPP, Zigbee Alliance, and WiMAX Forum at contributor and committee chair capacities. He is a past member of IETF Wireless and IPv6 Forum Technical Directorates. Alper is an author of numerous telecom-related standards and papers with 16 granted and several pending patents.

About the Company

Actility connects the industrial internet of things. Our IoT connectivity platform, tools, and fast-growing ecosystem enable our customers to create IoT solutions that transform business, industries and processes. The ThingPark LPWA platform connects sensors gathering data to cloud applications on any scale, from global or national networks to secure on-campus enterprise solutions, managing devices, data flows and monetization. Our value-added applications and business services enable roaming, device software update, geolocation and smart grid. Actility is at the heart of a thriving customer ecosystem, connecting solutions partners, supporting developers and device makers preparing their LPWA product for market, and providing an e-commerce Marketplace offering global distribution to solution providers. Actility co-founded the LoRa Alliance and continues to pioneer LPWA networking technology.

https://www.actility.com/

6.2 Firmware Updates Using LoRaWAN

By Jan Jongboom, Principal Developer Evangelist, Arm

Typical LoRaWAN deployments target more than 10 years of lifetime, but 10 years is a long time. Over a deployment's length devices might get repurposed, requirements change, standards evolve, and vulnerabilities emerge. Thus, firmware updates are essential for large-scale deployment of connected devices. Security patches protect customer and business data, and new functionality, optimizations, and specialization extend device lifetimes.

But how do we send the new firmware to our devices? A naive way of implementing firmware updates would be to split a firmware update into many packets, delivering these packets in the normal downlink windows. After every transmission a device opens a receive window, and we can deliver data back to the device. Unfortunately this implementation has a lot of downsides: Figure 1.

1. A firmware update might consist of 100 packets. If a device sends one message per hour, it will take well over four days to send an update.
2. There is no guaranteed quality of service (QoS) on the network, so devices need to indicate to the network which packets they have received, thus increasing network traffic.
3. LoRaWAN gateways have a limited downlink capacity. Sending large amounts of data from the network back to the device will exhaust this capacity and so deteriorate network quality.

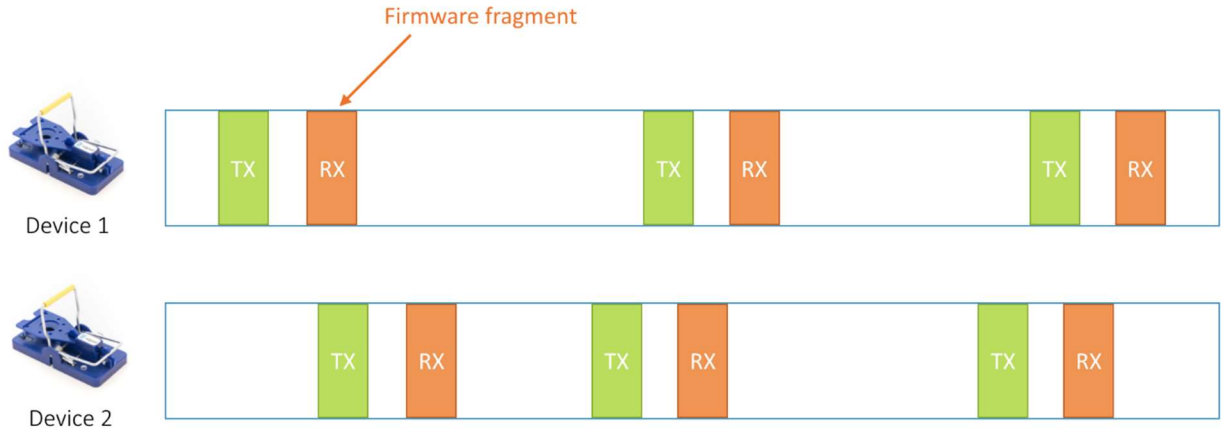

Figure 1: A naive approach, where firmware fragments are delivered to individual devices through their normal RX windows.

It's clear that this method does not scale. A much better method would be to send the firmware to all devices at the same time, so the gateway only needs to transmit twice. If the fragments are also sent in quick succession, without any device transmission in between, this approach is also much more energy efficient, as receiving consumes significantly less power than transmitting. Figure 2.

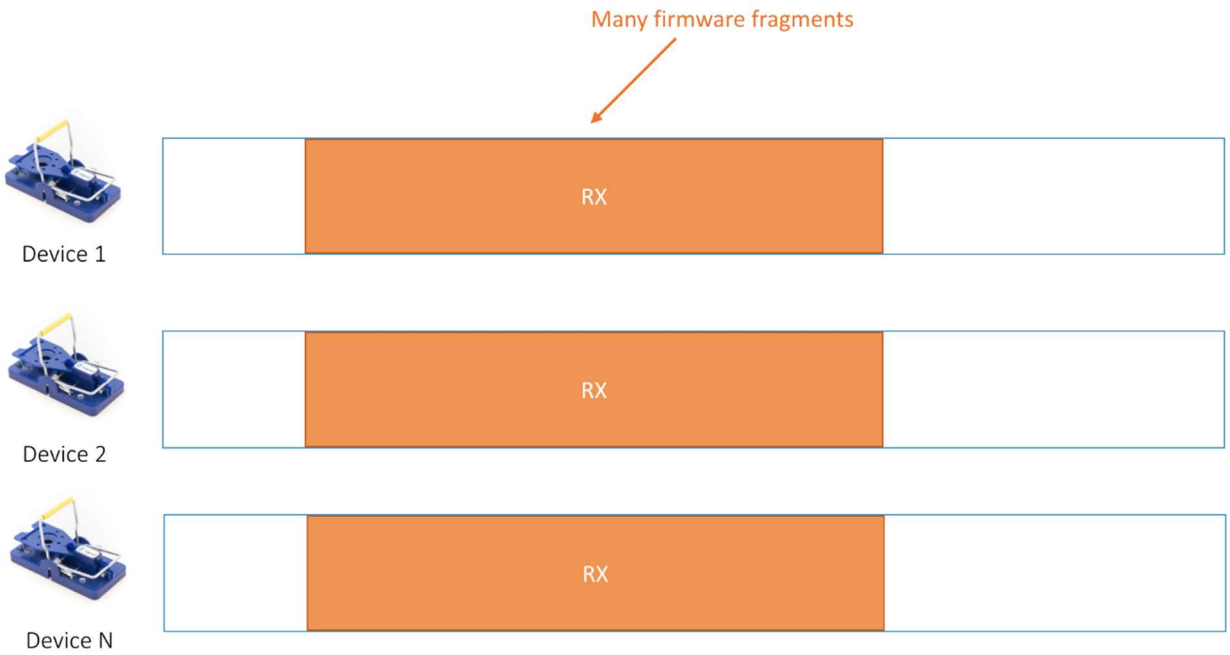

Figure 2: A better approach, where firmware fragments are delivered to all devices at the same time, and where no transmissions from the device are required.

However, as explained earlier, there is no guarantee that devices will actually receive every packet. It might be that packets are dropped. One way to determine if devices have received every packet would be to have devices request missed packets from the network, but this will lead to congestion on the network when lots of packets are missed, or when updating lots of devices. In addition, this approach requires the devices to send messages out, chipping away at battery life.

To mitigate the problems noted above we can use a technique called forward error correction. This is a technique to fix errors in data transmission over an unreliable communication channel, and we can use it to reconstruct missed packets. After sending the initial firmware we start sending correction packets. Devices that have missed packets will keep listening for these packets and reconstruct any missed packets using the correction packets. Once the device has reconstructed the full image it can stop listening. This adds some overhead, but it's much better than devices requesting specific packets. Figure 3.

Figure 3: The network delivers the regular packets first, then sends correction packets. Devices that have missed packets can recover these packets through the generic correction packets, rather than request a re-transmission.

New Specifications

To facilitate these techniques the LoRa Alliance has created three new standards, which were ratified in October 2018. These are:

- LoRaWAN Application Layer Clock Synchronization Specification.
- LoRaWAN Fragmented Data Block Transport Specification.
- LoRaWAN Remote Multicast Setup Specification.

LoRaWAN Application Layer Clock Synchronization Specification

LoRaWAN devices do not require an accurate clock, as they can transmit at any point. This is fine during normal operation, but for a firmware update we need to wake every device up at the same point. The clock synchronization specification introduces commands to tell devices the current actual time. Devices send a clock synchronization request to the network, containing the current value of some clock source that's present on the device. This can be the value of a Real-Time Clock (RTC), or the number of seconds since the device booted.

The network does know the accurate time at which the device sent this command, as the network does have a clock, and the network can calculate (based on the spreading factor) how long it took to send the message. It then compares this with the value of the clock source on the device and sends back a correction value. The device applies this correction value to its internal clock and now knows the actual time. If the clock drifts a few seconds over time, this is not a problem, as we can start the firmware update process a bit late to adjust for this clock drift.

LoRaWAN Fragmented Data Block Transport Specification

To know when the firmware image is complete the device needs to know how many packets are coming. Before the update the network thus sends a message that a fragmentation session will start at some point. This message contains the fragment size, the number of fragments, padding (as an image might not be aligned to the fragment size), and the forward error correction algorithm.

The fragmentation size is dependent on the spreading factor that will be used to send the image, as the maximum LoRaWAN packet size depends on this spreading factor. For example, SF7 / 125 kHz in Europe supports 222 bytes per package, but SF12 / 125 kHz supports only 51 bytes. You also need to subtract the size of the LoRaWAN header and the size of the fragmentation header (3 bytes). Note that this specification is not specific to firmware updates, but could be used to send payloads that do not fit in a single LoRa frame as well.

The forward error correction algorithm that's mostly used (and which is also used in interoperability testing within the LoRa Alliance), is Low-Density Parity Check coding (LDPC), but you are free to use a different algorithm.

LoRaWAN Remote Multicast Setup Specification

All communication over LoRaWAN is end-to-end encrypted, and this is also the case when sending the firmware to all devices at the same time (multicast). This does require however that all devices use the same device address and session keys during the firmware update. These multicast keys are sent to the device over a normal downlink message (encrypted with the device's session keys). After all devices have received the keys the multicast session can be started.

To start a multicast session the network sends a message containing the time that the session starts (the devices know the time through the clock synchronization spec), the frequency that is used, the spreading factor and bandwidth for the session, and a timeout value. When the multicast session starts all devices then switch to Class C (continuous listening mode), with the keys and other parameters set to exactly the same values.

The network can now send a single message, which can be decrypted by every device in the multicast session. Note that the multicast session is vulnerable to man-in-the-middle attacks, as every device in the multicast group has the same keys and could send out rogue messages. You should thus always verify that the final firmware image came from a trusted party, for instance via a public / private key pair (see later in this article).

Tips for a Successful Update

While these specifications make firmware updates possible, these operations will still have a significant effect on your network. Here are some considerations when doing a firmware update:

1. Don't update the full application. Implement a delta update algorithm to only send the changed parts of the application. This can reduce the size of an update by 90%.
2. If you're in a region with plenty of spectrum available (such as the U.S.), use a dedicated frequency band for firmware updates.
3. Send as fast as possible. If 98% of your devices can be reached on SF9 and 2% on SF11, it is probably better to only update the devices on SF9 and use a different mechanism to update the other devices. Lower data rates require exponentially more time and support smaller fragment sizes.
4. If you're in a region with duty cycle constraints (such as Europe) consider the available downlink capacity on gateways before scheduling an update.
5. Duty cycle is measured per hour, so bursting for two minutes on 869.525 MHz (10% duty cycle in Europe) is actually fine. Make sure you keep some spare capacity for the RX2 windows.

Security in a Multicast Session

When instructing multiple devices to join a temporary multicast session where all the devices share the same session keys, there is a potential security risk when one of the devices gets compromised: packet injection. An attacker with the multicast session keys can send packets as if they came from the server. Theoretically it could also be that your network provider gets compromised, and an attacker initiates a multicast session with new rogue firmware.

Given how critical firmware update service is (if an attacker can send its own firmware you might never get the device back), you should thus add an extra layer of security. A proven way to verify that a package was sent by a trusted party is to use asymmetric cryptography. A public / private key pair is created by the manufacturer, and the key is loaded into the device's Read-only Memory (ROM). When an update is created the update is signed with the private key, and the signature is included in the update.

When the device receives the update it verifies the signature against the file it received and the public key held in ROM. Only if this verification passes does it proceed to update the device. A way of doing this efficiently on embedded devices is to use ECDSA-SHA256. Here the SHA256 hash of the new firmware is signed using elliptic curve cryptography. In addition, you can add some extra sanity checks, such as including a manufacturer ID and a device type ID to ensure that the right firmware was sent to the right device.

For more information: The IETF SUIT working group (https://datatracker.ietf.org/wg/suit/about/) is actively working on specifications for secure firmware updates of embedded devices.
Note that it's of the utmost importance to keep the private key private. Keep the key on an air-gapped computer and limit the number of people that have access to the keys. Figure 4.

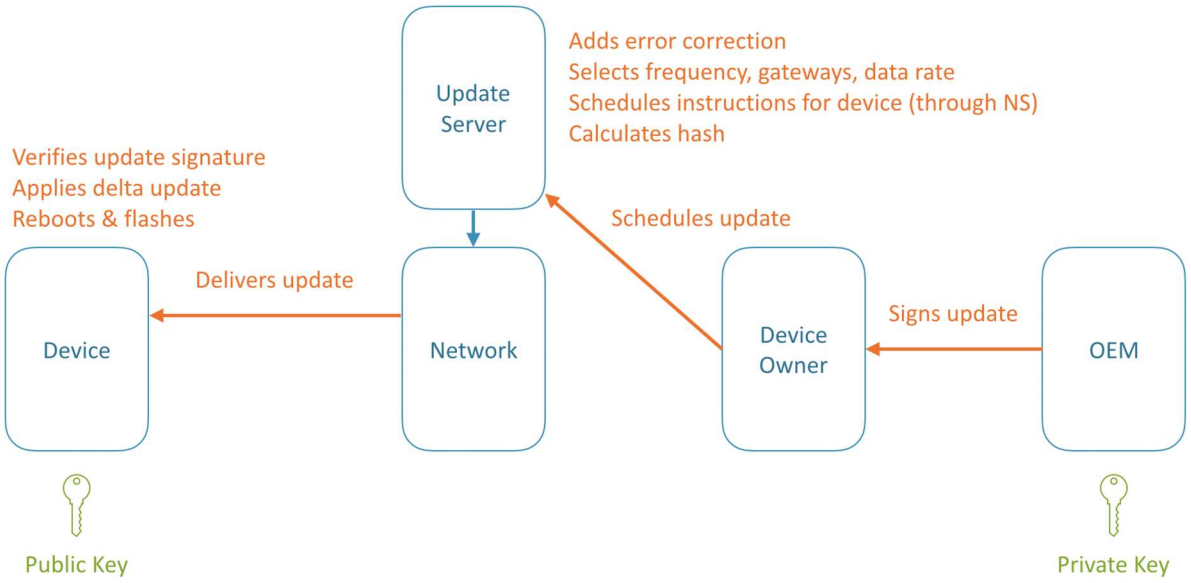

Figure 4: A fully secure update flow where the update is signed by a private key, which is held by the OEM or manufacturer. The network is required to deliver the update but cannot manipulate the content of the update. The device verifies the cryptographic signature through the public key before applying the update.

How Long does it Take?

All the above was not just theoretical. Firmware updates are actually feasible, even when dealing with a duty cycle. If we take an update with the following properties:

- Data rate: SF9 125 kHz (in Europe).
- Delta update size of 8K.
- 5% packet loss.

At this data rate the maximum fragment size is 110 bytes, which takes 595ms to transmit. This gives 82 packets required for an update (10% redundancy packets used) for a total of 48 seconds of transmission required. This 48 seconds is enough to update all devices in the vicinity of the gateway. In Europe there is a 10% duty cycle on 869.525 MHz, which is six minutes per hour. So by using 48 seconds for our update there is still plenty of capacity left.
(Note: You might want to add some pause between packets to ensure RX2 can still be used on your network.)

In the U.S. FCC regulation imposes a maximum dwell time of 400ms on uplinks between packets for a total of 81 seconds; but you can use a dedicated update frequency due to the wide amount of spectrum available.

Device Requirements

There are also some considerations to make when designing your device to be ready for firmware updates. The most important one is that you need a place to store firmware fragments, typically in external flash. The amount of storage required depends on the maximum size of your firmware and whether you want to support delta updates. Delta updates require more space, as both the incoming delta patch, the old firmware, and the patched (outcome) version need to be stored.

In addition you might want to consider storing a known good version of your application, so you have a way of restoring the device if a firmware update goes wrong. Thus, for a device with 256K internal flash we'd recommend adding 1M of external flash. If you don't want to support delta updates, and if you have enough flash on your MCU, you could store the update in internal flash, too. For example, you could partition a device with 512K internal flash into two banks: one for the current firmware, and one to download new firmware into.

You also need a bootloader in place. A bootloader is a separate application that runs before your main application, and it's responsible for flashing your device with new firmware. It's wise to keep the bootloader as small as possible, as updating the bootloader is often not possible without physical access to the device. To establish a trust relationship between your application (which receives the new firmware) and the bootloader set up a Root of Trust (RoT). A RoT is a secret key, which the application uses to prove that it wrote the new firmware. This technique prevents a hacker from physically connecting to your flash and writing a new firmware image directly, as it will not be signed with the RoT.

Last, take into account that proper cryptographic verification of the new firmware is resource intense and requires some Random Access Memory (RAM). Verifying the authenticity of firmware through ECDSA/SHA256 could take up to 10 seconds and use 6K of RAM on a small device. Thus, make sure your device has enough RAM to spare, or implement memory pressure events, so the rest of the application frees memory when the verification happens.

The time delay can be offset by using an external crypto accelerator which can do the verification. This also has some other benefits, such as keeping keys always in a tamper-proof component, but adds cost.

Conclusion

Firmware updates are crucial for any long-term deployment. They can add functionality, fix vulnerabilities, and adapt deployments to changing regulations. With the new multicast, data fragmentation, and clock synchronization specifications, it's now possible to efficiently distribute new firmware to many devices at the same time.

If you want a quick start in development, device side reference designs are available from Semtech and Arm (https://github.com/armmbed/mbed-os-example-lorawan-fuota). The Semtech stack implements the three specifications, and can be used for interop testing. The Arm stack is a full firmware update client, implementing the specifications, bootloader, and delta updates. Network operators will start supporting the specifications in 2019, so refer to your network operator to find out when you can start updating your LoRaWAN devices in the field!

About the Author

Jan Jongboom is an embedded engineer and Developer Evangelist for the Internet of Things at Arm, always looking for ways to connect more devices to the internet. He has shipped devices, worked on the latest network tech, climbed upon buildings to install gateways, and there's a monument in San Francisco with his name on it. He has been working with LoRaWAN since 2015 and is a member of the LoRa Alliance Technical Committee.

About the Company

Arm technology is at the heart of a computing and connectivity revolution that is transforming the way people live and businesses operate. Our advanced, energy-efficient processor designs have enabled intelligent computing in more than 145 billion chips and our technologies now securely power products from the sensor to the smartphone and the supercomputer. In combination with our IoT device, connectivity and data management platform, we are also enabling customers with powerful and actionable business insights that are generating new value from their connected devices and data. Together with 1,000+ technology partners we are at the forefront of designing, securing and managing all areas of compute from the chip to the cloud.

https://www.arm.com/

6.3 Accelerating IoT end node design with LoRaWAN

By: Ramya Kota, Product Line Manager, Microchip

For decades, wireless developers have faced the dilemma of choosing between longer range connectivity or lower power consumption. LoRa® technology has emerged as a compelling mix of long-range connectivity, low power consumption, and secure data transmission. It is now the leading option for battery-powered end devices that need long range connectivity. With the growing Internet of Things (IoT), LoRa wireless technology solutions address increasing demands on end devices while offering low infrastructure cost for volume deployment.

This article will introduce the four main elements of LoRaWAN network architecture and discuss LoRaWAN end nodes in detail, including their classifications and various options available for designers to quickly deploy LoRaWAN end nodes. It will also discuss some of the challenges faced by designers while developing LoRaWAN end nodes and how to overcome those challenges by choosing the right System-on-Chip (SoC)/System in Package (SiP) or module.

LoRaWAN Network Architecture

LoRa is a wireless modulation technique or physical layer that allows low-power end devices to communicate over a long range. LoRaWAN is a wireless networking protocol that acts as a media access control (MAC) layer and is implemented on top of the LoRa physical layer. The LoRaWAN specification details the communication protocol and network architecture and is meant to provide secure communication for end devices and interoperability within the network. Figure 1.

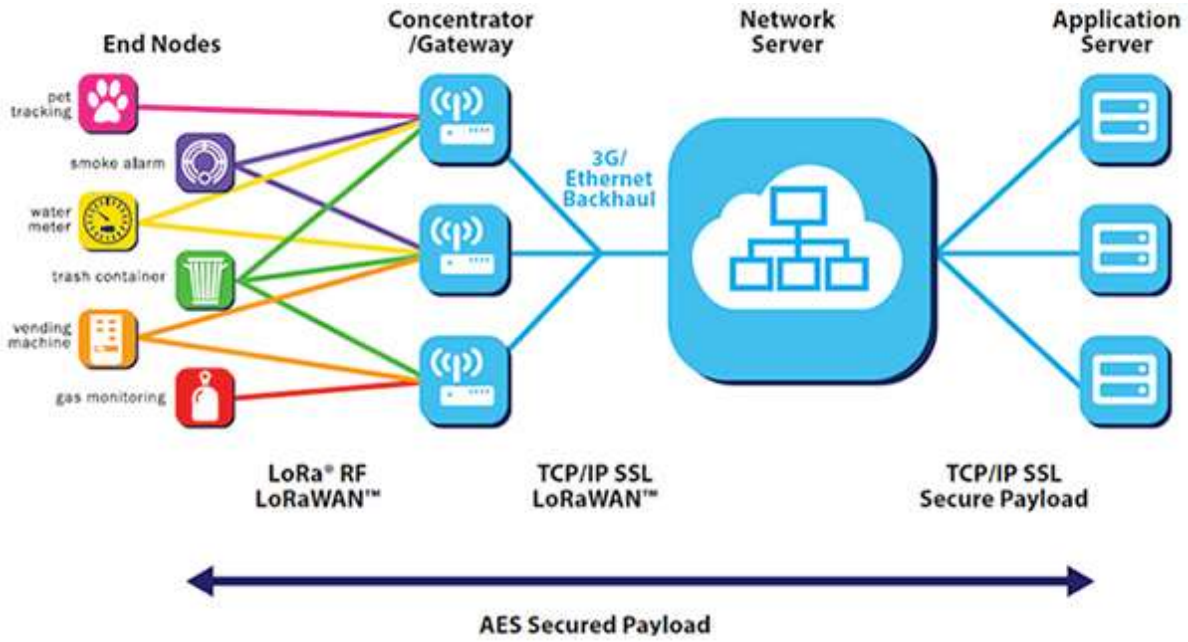

Figure 1. The four elements that comprise the LoRa network (Image source: LoRa Alliance)

The LoRa network has four elements:

1. **End nodes** are elements of the LoRaWAN ecosystem that gather sensor data and transmit/receive the data. These are generally remotely connected and are battery-operated.
2. **Gateway** is a transparent bridge between the end nodes and network server. Typically, end nodes use LoRaWAN to connect to the Gateway, while the Gateway uses high-bandwidth networks such as Wi-Fi®, Ethernet, or Cellular to connect to the networks.
3. **Network Server** connects to multiple gateways. It gathers data from the gateways and filters out duplicate messages, decides which gateway should respond to end node messages, and adapts data rates to extend the battery life of end nodes.
4. **Application Server** collects data from and controls the actions of the end nodes.

Let's take a closer look at the LoRaWAN end node architecture and some of the key factors to consider when choosing the best devices for LoRaWAN end nodes.

LoRaWAN End Node Architecture

End nodes are simple objects such as sensors and actuators. Typically, these are the "things" in the Internet of Things (IoT). In the LoRaWAN ecosystem, an end node communicates to the network server through one or many gateways.

The LoRaWAN specifications define three classes of end node devices. All three device classes communicate bidirectionally and can initiate an uplink to the server via the gateway. They differ regarding when the end node listens to the network server.

Class A

A LoRaWAN Class A end node offers the lowest power consumption and is ideal for battery-operated applications. A Class A end node starts each transaction with an uplink transmission, which is then followed by two downlink receive windows. The network server sends the downlink message after receiving the uplink. At the end of downlink message, the end node enters sleep mode, thereby saving power. This allows Class A end nodes to consume the least amount of power and provide extensive battery life. All LoRaWAN end nodes support Class A by default.

Class B

In Class B, the end node reduces the downlink latency by opening periodic downlink receive windows, unlike in Class A where the downlink is non-deterministic. The periodicity of the downlink windows is maintained by synchronizing the clocks of the end node and the network server. For the synchronization, the network server commands the gateways to send a beacon at regular intervals.

Class C

A LoRaWAN Class C end node continuously opens the receive windows with the exception of the uplink period. Class C end nodes reduce the latency considerably but also increase the power consumption significantly.

How to Simplify LoRa End Node Development

Depending on an application's development time, target costs, power consumption, and the developer's radio frequency (RF) expertise, there are several options available to build LoRa end nodes. They can primarily be classified in two categories: System-on-Chips (SoCs) or System in Packages (SiPs) and modules.

LoRa SoC and SiP Solutions

A LoRa SoC or SiP includes a microcontroller (MCU) and LoRa transceiver in a single package. The biggest advantages of SoC- or SiP-based end node designs are their abilities to achieve compact end node size and lower system cost.

While these solutions provide significant flexibility, size, and cost benefits, end node developers also have challenges to consider when choosing SoCs/SiPs.

The most common challenges when designing this end node architecture include:

1. **RF design expertise**

   When using SoCs/SiPs, the end node developer is responsible for the entire RF design, including schematics, Bill of Materials (BOM), Printed Circuit Board (PCB) layout, antenna tuning, and other RF hardware considerations. The RF design not only requires in-depth RF expertise, but also adds significant development time for designers. To overcome this challenge, it is important to choose SoCs/SiPs that are supported by regulatory-certified reference designs and detailed chip-down design packages. For example, Microchip's SAM R34/35 LoRa SiP devices provide a detailed chip-down design package that includes schematics, BOM, Gerber files and certification, and RF testing tools. Using thoroughly tested and certified chip-down designs helps reduce risk and development time significantly.

2. **Regulatory compliance**

   Another challenging aspect of using SiPs/SoCs is regulatory and standards compliance. LoRa/sub-GHz radios typically operate in the industrial, scientific and medical (ISM) license-free band. Frequencies vary depending on the region, making it challenging for hardware and software designers to design products that work across geographies. Diligent care must be taken to design a fully compliant solution while keeping the BOM costs minimal. Furthermore, repeated certification testing increases the certification costs as well as the

development time for end nodes. To avoid this issue, it is essential to choose SiP/SoC solutions that provide certified reference designs along with all the required RF and certification testing tools. Microchip's SAM R34, for example, comes with certified reference designs and allows developers to use a single part variant across geographies with support for worldwide LoRaWAN operation from 862 to 1020 MHz. Knowing that the SiP/SoC-based reference design is regulatory compliant in a region of choice gives designers ease and significantly reduces the design risk.

3. Proven software

Generally, LoRa modules integrate the whole LoRaWAN stack inside the module, and the end node developer only needs to implement the initialization and communication to the module. With LoRa SoCs/SiPs, the manufacturer must provide the stack, or the developer must develop their own stack if no stack is provided. To minimize software development time, it is recommended to choose a SiP/SoC that is supported by a manufacturer's LoRaWAN stack. Using proven LoRaWAN stacks from manufacturers ensures interoperability of end nodes with major LoRaWAN networks and gateways, enabling end nodes to work across different networks with reduced risk.

4. Power consumption

Apart from long-range connectivity, one of the key features for most LoRaWAN applications is low power. Most LoRaWAN end nodes are battery-operated, remote sensor applications that only wake up when needed and sleep most of the time. These applications ideally require minimal energy and achieve prolonged battery lifetime of up to 10 years, thus minimizing battery replacement costs. End node designers should carefully consider the tradeoffs among costs, features, and power when selecting a SoC for their application. Choosing a LoRa SoC that has various power domains and ultra-low-power features would be ideal for preserving and extending the battery life of remote, connected end devices.

SAM R34/35 devices provide sleep modes as low as 790 nA to significantly reduce power consumption and extend battery life in end applications.

Overcoming Design Challenges with LoRa SiPs/SoCs

To overcome these challenges, developers can choose LoRa SiP/SoCs that are supported by certified reference designs, hardware RF guidelines, and interoperable software. By doing so, developers can reduce development time while enjoying the benefits of small size and system costs.
For example, the SAM R34/35 comes with certified reference designs, proven LoRaWAN software, ultra-low-power consumption, and compact packages to help developers overcome traditional challenges and accelerate time to market. Figure 2.

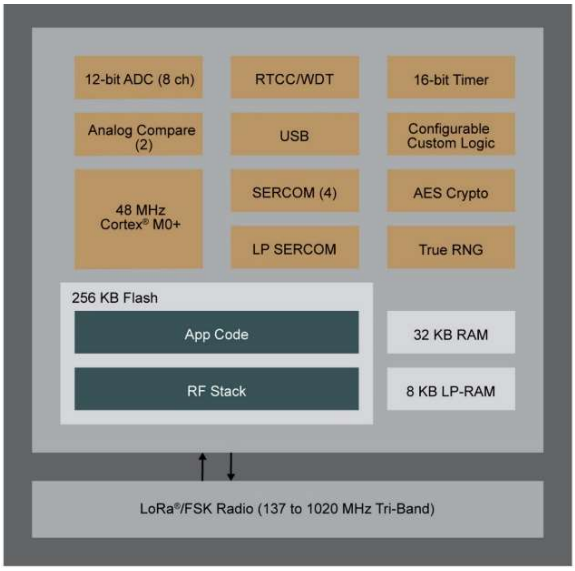

Figure 2. Block diagram for Microchip's SAM 34/35 family

LoRa Modules

Designed specifically to simplify development of LoRaWAN end devices, LoRa modules contain all the required radio components along with the LoRaWAN stack and RF circuitry. The main advantage of LoRa modules is that the RF hardware development is implemented by the module manufacturer. Additionally, the entire matching circuitry is inside the module provided by the module manufacturer, saving significant development time for end node designers. There are two primary types of LoRa modules available.

1. **Standalone modules** are the LoRa modules that do not require any external host processor. These modules not only include all the required radio components but also include enough memory to run the application code along with the LoRaWAN stack. The standalone modules are ideal for simple, new LoRa applications that require very quick time to market

2. **Coprocessor modules or modems** contain all the radio components including stack and RF circuitry but require an external microcontroller to manage them. With reference designs and guidelines, these modules can easily be connected to an external antenna. The communication to the host is managed through attention (AT) commands to configure the module or to send data via Universal Asynchronous Receiver-Transmitter (UART), USB or System Packet Interface (SPI). The application code is external to the module, giving designers extensive flexibility in selecting a host MCU of their choice or using an existing host MCU that is familiar to them. These modules are ideal for existing and new applications that require designing LoRaWAN applications without much RF expertise.

While LoRa modules accelerate designs with simplicity and ease of use, they can be expensive and bulky for cost- and size-sensitive LoRa end node applications. Choosing a low-cost, compact LoRa SiP or SoC is ideal for such end node applications, as they address lower size and cost requirements. Ideally, LoRa modules are a perfect fit for applications that require less RF expertise and faster time to market. Selecting an easy-to-use module that can communicate to the host MCU via simple UART saves significant development time and provides the fastest way to market.
A simple example of one such solution is shown in Figure 3, where the certified LoRa module provides a drop-in solution for LoRaWAN connectivity with its onboard command processor, LoRaWAN stack, radio transceiver, and UART connectivity. Figure 3.

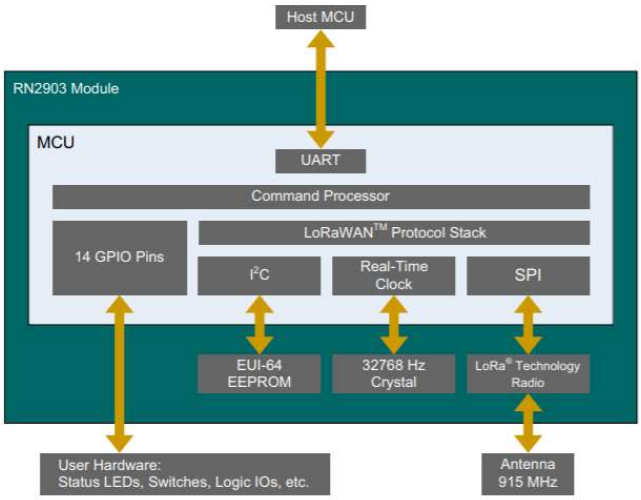

Figure 3. Block diagram for Microchip's RN2903 LoRa module

Conclusion

With the growing Internet of Things, LoRa technology satisfies the essential IoT need to enable low-power devices to communicate over long ranges. Compliant RF designs, interoperable LoRaWAN software, and low power consumption are factors to consider when developing LoRaWAN end nodes. Selecting the right wireless SoC/SiP or module while considering the tradeoffs between lower costs and faster time to market are the keys to a successful LoRaWAN design.

About the Author

Ramya Kota is a product line manager for Microchip specializing in LoRaWAN products. Her responsibilities include new product definition and marketing for Microchip's wireless products. Prior to joining Microchip, Ms. Kota served as a product line manager for Silicon Labs, managing their Interface and Bridge Portfolio. She holds a master's degree in computer engineering from the State University of New York at Stony Brook.

About the Company

Microchip Technology Inc. is a leading semiconductor supplier of smart, connected and secure embedded control solutions. Its easy-to-use development tools and comprehensive product portfolio enable customers to create optimal designs which reduce risk while lowering total system cost and time to market. The company's solutions serve more than 125,000 customers across the industrial, automotive, consumer, aerospace and defense, communications and computing markets. Headquartered in Chandler, Arizona, Microchip offers outstanding technical support along with dependable delivery and quality.

https://www.microchip.com/

6.4 Intelligence at the Edge

By Sara Brown, Vice President of Marketing, MultiTech

The Internet of Things (IoT) has set the IT world on fire with innovative technologies and astonishing growth levels. Forbes has even forecast this booming trend to continue with the IoT market predicted to reach $267 billion in 2020[1]. With this growth comes an incredible amount of data being collected, which will accelerate in coming years. An IBM Marketing Cloud study claims 90% of the data in the world today has been created in the last two years alone, at 2.5 quintillion bytes of data a day! At this rate, it won't be long before we overwhelm data centers and burden the electrical grid that powers and cools them.

How can we meet the needs of this exponential growth in data with more and more smart-connected sensors and assets being deployed daily? Innovations in machine-to-machine (M2M) and IoT networks are necessary to alleviate the costly inefficiencies of data handling centrally in the cloud. IoT networks must embrace a decentralized architecture where processing of raw data begins at the network edge.

The Opportunity for Intelligence at the Edge

Enterprises now recognize that the ever-increasing flow of data must be managed more efficiently in order to optimize information use, reduce costs, and improve business performance.

Ineffective data management systems create situations where laborious and costly data processing must be performed centrally in the cloud or at an enterprise's data center, thus increasing operating costs. Examples include:

- Backhauling data, often over costly cellular networks
- Processing raw field-bus data from industrial, often proprietary assets, into a format that can be exchanged more efficiently by Internet-based platforms
- Data storage that is accessible when needed to generate actions and automate business processes, such as, upon what conditions to accept delivery of a shipment of frozen goods or perform a truck roll to deliver or inspect an asset

[1] https://www.forbes.com/sites/louiscolumbus/2017/01/29/internet-of-things-market-to-reach-267b-by-2020/#534e5ad609bd

Whether virtualized or physically located in an enterprises data center, ongoing operating expenses are generally billed by data flow in and/or out of the cloud, or based on the number of cores or racks of equipment purchased to perform calculations and store the results. Often raw data being sent upstream, such as minute changes in temperature or humidity, has little or no value. Processing and storing this raw data for a period of time only to eventually delete it without its ever having never been used in a meaningful business process, is costly and wasteful. Capturing and immediately sifting information at the network edge and making decisions locally while securely transferring only actionable data and exceptions to the cloud is paramount to business success. Moreover, this method enables certain actions to be taken in situ, reducing latency for mission-critical, real-time applications.

As artificial intelligence and machine learning matures, more and more of this decision making can and should reside at the edge, requiring increased processing power and logic to be built into edge devices. Finally, the ever-changing data security landscape requires adequate processing overhead to execute security algorithms, store and encrypt keys, authenticate messages, scan updates, etc., in order to protect critical infrastructure.

To act on this data with timely and informed business decisions, intelligent data management systems at the edge of the network are critical. Utilizing efficient data management at the edge of the network eliminates the need for time-consuming data back-hauls and the parsing required to prepare data for business decisions.

The Case for Unlicensed Networks

A new class of unlicensed, low-power, wide-area networks (LPWANs) can free industrial applications from the consumer-driven cycles of the public cellular networks. Instead, LPWANs offer the stability of public or private networks designed and built specifically for machines. These networks extend battery life and range and provide connectivity for the large majority of connected device use cases. LoRaWAN is an LPWAN protocol intended for wireless, remotely located, often battery-operated “things” in local, regional, national or global networks. LoRaWAN provides secure bi-directional communication, mobility and localization services, and seamless interoperability among smart things. Better yet, thanks to its low power consumption, a cost of pennies on the dollar compared to alternative technologies, and the ability for enterprises to deploy their own private networks for added security, LoRaWAN has the opportunity to facilitate the long-predicted huge IoT growth with just-right connectivity for sensor-based IoT applications.

Intelligence at the Edge of LoRaWAN Networks

LoRaWAN opens the door to a whole new class of connected devices. Incorporating programming and processing capabilities at the edge of LoRaWAN networks enables optimal distribution of decision-making right on time at the very source of the actionable data. Now it’s possible to affordably connect the unconnected, while simultaneously limiting potential strain on networks and data centers.

For example, to achieve success in agricultural settings, farmers need to understand the frequent fluctuations of temperature, humidity, and soil moisture. Moreover, they may need to take corrective action, should the balance of factors be outside the plants' tolerance. Armed with intelligent LoRaWAN sensors and smart, programmable gateways, farmers gain up-to-the-minute data about the growing conditions on their farms, empowering them to automate corrective action right at the site, without having to pay for expensive data transmission and cloud storage or tolerate the latency inherent in transmitting that information to the cloud and back.

Similarly, facilities managers at hotels, hospitals, and schools have a great deal to monitor— from access control to room occupancy to water temperature— even how many paper towels are left in a given dispenser. By equipping their facilities with intelligent LoRaWAN edge infrastructure, they can improve sanitation, reduce costs, and automate corrective action without the need for another server room.

MultiTech Solutions

MultiTech provides solutions solve the complex problem of data management by assisting in decentralizing IoT architecture by collecting and interpreting data at the network edge. Distributed intelligence delivers critical data processing and facilitate immediate business decisions. With the capability of analyzing data in real time, MultiTech enables you to make use of data when it is most important minimizing data processing and storage costs. Additionally, software upgrades and system maintenance can be done remotely.

MultiTech has a suite of intelligent, programmable devices across our portfolio to support edge processing and optimize your data management. The MultiConnect® xDot™ and MultiConnect® mDot™ are low power RF modules that feature a long range of up to 10 miles line-of-sight or three miles in buildings, as well as low bit rate M2M data connectivity to remote sensors and industrial equipment. The xDots and mDots reduce complexity and deliver intelligence to the very edge of your LoRaWAN network. Pairing these smart modules with a MultiConnect Conduit smart programmable gateway enables even more processing and decision-making process without further burdening either the network or the data center.

For a growing number of businesses, data collection has surpassed data processing. The opportunities missed by underpowered or nonexistent data management systems at the network edge can have substantial business consequences. From U.S.-based manufacturing facilities, MultiTech delivers the industry's leading data management products to deliver high performance and intelligent data processing. It's no wonder you can find more than 25 million MultiTech products in service around the world today! Learn more about our intelligent edge infrastructure by visiting us online at MultiTech.com.

About the Author

As Vice President of Marketing Sara Brown is responsible for product marketing, as well as internal and external marketing and communications. She brings 20 years of technology marketing experience, with more than ten years focused on M2M and IoT. In addition to her experience as a senior creative at some of the world's premier advertising agencies, Sara has worked for IBM, Wavecom/Sierra Wireless and Telit Wireless Solutions, and consulted to global technology companies. She has served as Marketing Co-Chair for the LoRa Alliance and is currently Vice Chair of the IoT M2M Council.

About the Company

MultiTech designs, develops and manufactures data communications equipment for the industrial Internet of Things — connecting physical assets to business processes to deliver enhanced value. MultiTech is committed to quality, service and technology innovation.

https://www.multitech.com/

6.5 Microshare Smart Facilities Management Solutions on LoRaWAN®

The promise of LoRaWAN will be realized with IoT solutions that support your business initiatives and grow and change as your business grows and changes.

Microshare offers Smart Facilities Management accelerators in the form of starter kits that enable you to deploy a working IoT solution on a LoRaWAN network in less than a day – without IT involvement.

By Charles Paumelle, Ellen Brockley and Tim Panagos, Microshare, Inc.

The explosive growth of IoT sensor and network technologies promises to change the way we live, work, and earn in ways we cannot yet imagine. Innovation is lowering the cost of ownership and increasing ease-of-use to enable ubiquitous IoT adoption across every industry, home, and office around the world. The LoRaWAN infrastructure is available, evolving, and fast becoming the de facto standard for IoT across the board.

Now we need IoT solutions that deliver real-time and historical data – the kind that fuels insights for better decision-making. IoT solutions should be cost-effective and flexible enough to adopt new technology innovation as it emerges. They should collect and protect data from multiple sources. Business users should be able to consume and leverage IoT data with minimal training or technical expertise. Data should flow to the people who need it at the right time on their preferred devices with complete control and compliance.

In short, IoT solutions should *work out of the box* supporting your business with insight that helps you make better decisions and protects your data. IoT ecosystems should grow seamlessly and change with you – from your first installation to your ultimate smart facility. You don't want to be locked down with proprietary devices, data silos, and me-too apps. You want data to be available but secure.

There are many questions when embarking on your IoT journey. Where do you start? Which sensor technology is best for your use case? What expertise do you need to build a whole solution? How will you deploy? Will you have to train users and convince them to leverage new information to optimize your IoT value? How long is all of this going to take?

Definitively answering all of these questions up front wastes valuable time and precludes taking an early advantage to create value. That is, you leave a lot on the table if you insist on developing your long-range plan up front. You will be stalled by the inertia of inactivity. What you need is forward momentum. It is best to just get started with a platform that gives you the confidence that your IoT ecosystem will evolve with technology advancements, user acceptance, and your own knowledge and expertise.

To that end, Microshare offers a data sharing platform that supports any IoT solution from any industry. The platform is technology-agnostic, so you can plug and play with new devices and

other solutions as they emerge. To help you get started Microshare also provides solution accelerators in the form of starter kits so you can deploy a working IoT solution on a LoRaWAN network in a matter of days – without IT involvement.

Let's evaluate the framework with examples from facilities management.

Microshare Data Ownership Framework

Microshare is a simple tool to use. There is one RESTful API call to inject data (via POST or web socket), and there is one RESTful API call to retrieve data. Behind the scenes is a data lake that stores the introduced content that has been decorated with metadata. When you want the data back, you call for it. No matter what the format or source, you get your data back from the data lake. See Figure 1.

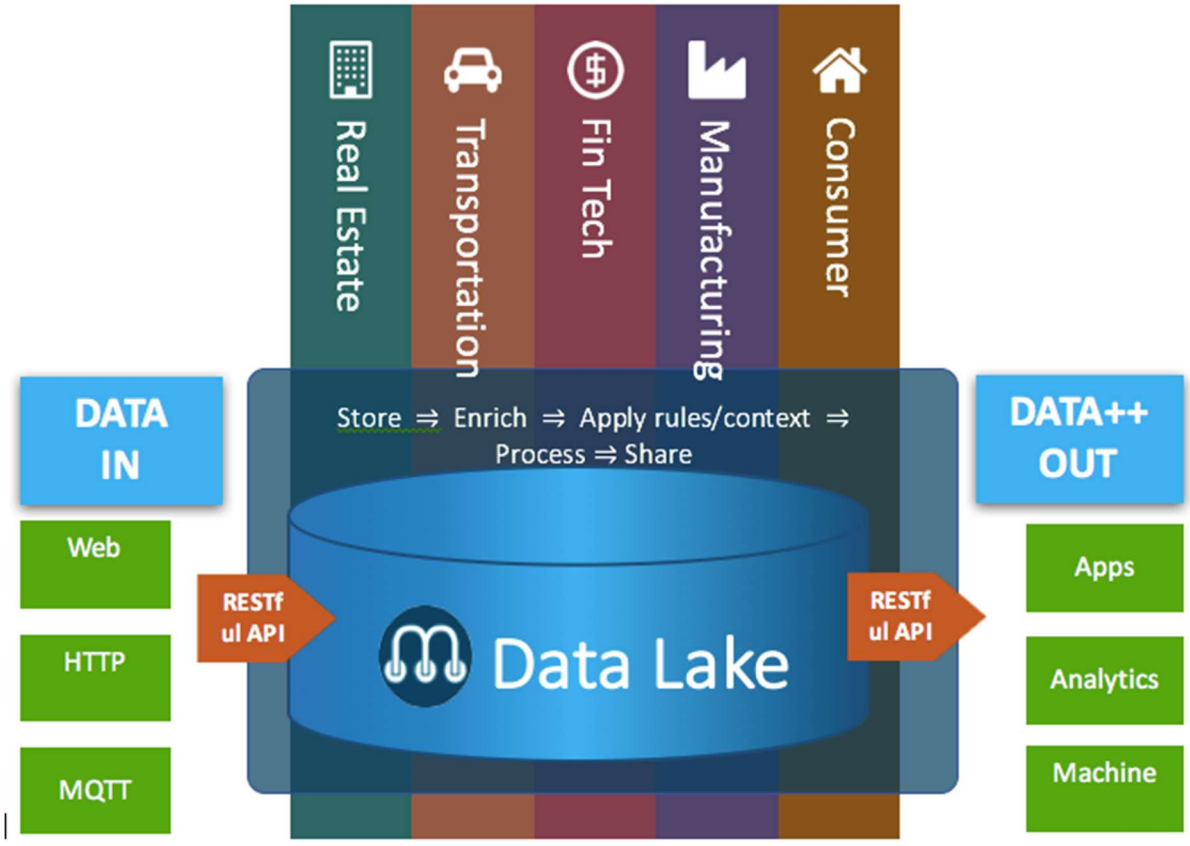

Figure 1. Microshare Data Ownership Framework

The magic of Microshare is the automatic application of data governance policies to determine which data may be retrieved in any given circumstance. Microshare uses three important elements to make this application possible:

1. Ownership metadata collect when an object is created;

2. Request metadata derives when an inbound request is made;
3. Rules that codify the conditions and context of access.

Data owners have the right to set the Rules that determine who, when, and how their stored data may be accessed by any Requestor. In simple terms, Rules determine what data is returned, to whom, in which circumstances. Inside of Microshare is a purpose-built Rules Engine that combines the context of ownership with the context of the request and evaluates your established policy to determine which 'rows' and which 'columns' of data are appropriate to return. Rules can show one set of data to party A and a different set to party B without the need to duplicate that data or perform costly offline transformations. See Figure 2.

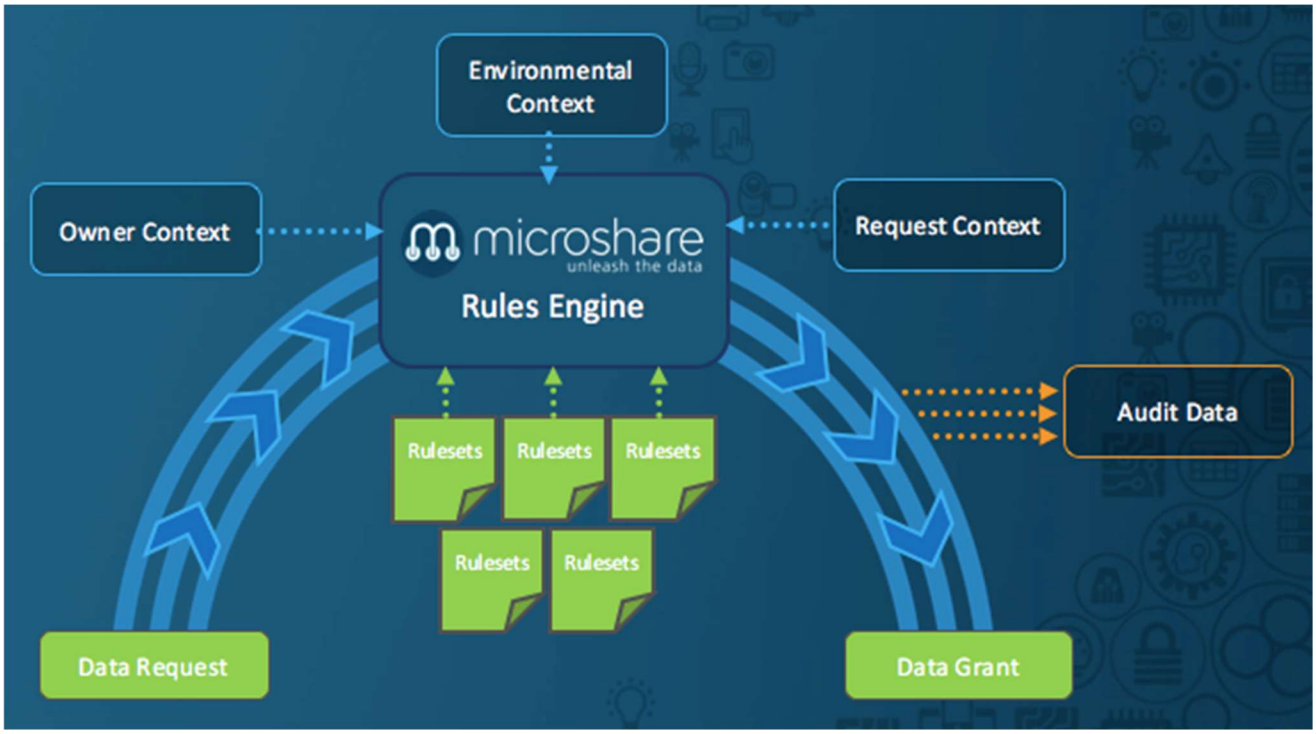

Figure 2. Microshare Rules Engine

By using the single Microshare GET API to interact with stored data, the end users, their application software, and the developers that support them need not be aware of the centralized governance decisions. That's because the data and the API remain constant even as the underlying policies change to follow evolving business needs. The whole system allows organizations to get started quickly with the assurance that their governance processes can continue to evolve without the need for costly rewrites or software crashes. Organizations using Microshare can act without fear of making irrevocable mistakes and focus on the needs of the business and the demands of the regulatory environment.

All requests, the Rules considered, and the important context used to evaluate the outcomes are logged durably to feed reporting, compliance, fraud, and analytics tools and dashboards. Running GET & POST activity through Microshare gives you a complete log of all data access regardless of source and destination.

That is the core value proposition of Microshare. The right data is made available to the right people. That is true from day one, but it is never set in stone. Microshare applies to data shared amongst users, across business units and geographies, and between organizations with equal ease.

Microshare Solutions – Smart Facilities Management

With the Microshare data sharing platform as the foundation, you can start building solutions that address your specific needs. A good place to start is commercial real estate in general and facilities management in particular. Operational improvements in cleaning and maintenance as well as space utilization not only save on costs but also improve efficiency and customer experience.

Microshare allows you to manage your entire operation with the power of real-time information about how people really use your space. With that knowledge you can make better decisions about resource allocation, space utilization, and environment quality that will reduce cost, optimize efficiency, and improve customer experience.

Microshare enables you to tailor IoT solutions with your choice of thousands of sensing devices to meet the specific needs of your business. Our starter kits get you up and running in a single day – without IT – so you can reduce costs, improve efficiency, and delight your customers all at the same time.

Starter kits include sensor and gateway hardware, application software in the cloud, and simple instructions to deploy your solution. See Figure 3.

1. Unpack the box.
2. Attach the sensors in your desired locations.
3. Name your locations on your sensor map and upload to the Microshare cloud.
4. Plug in the gateway.

Within 24 hours your dashboards will display your data – in real-time

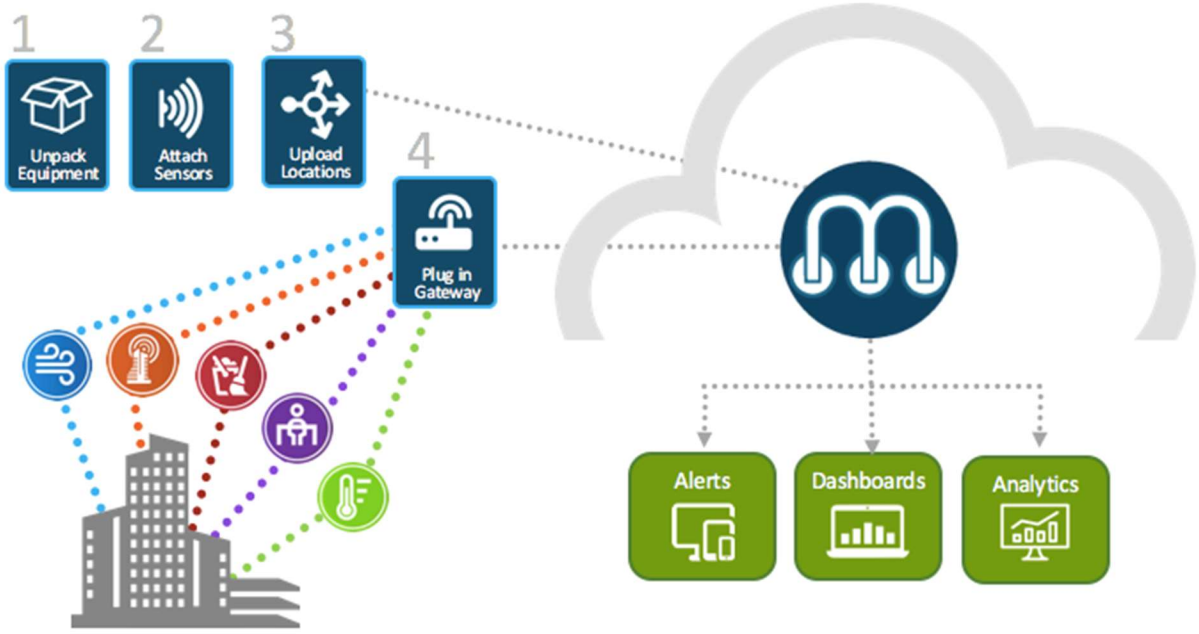

Figure 3. Microshare Starter Kits

Microshare offers your choice of three solutions to get started:

- Predictive Cleaning
- Desk / Meeting Room Occupancy
- Healthy Office

Data from each solution is presented in a specialized dashboard and can be combined with any other data source to provide a holistic view of your entire operation.

Predictive Cleaning

Most facilities managers spend significant budget cleaning areas that are already clean, emptying waste bins that are only partially full and responding to predictable crises because they don't have real-time insight about when to intervene. This approach is expensive and inefficient.

Data generated by actual use allows you to move from wasteful schedule-based cleaning services to more efficient just-in-time cleaning. For example, you could set a usage threshold on a space that would trigger an alert to a cleaner when that threshold was met.

With insight into usage patterns you can transition from simply reacting to crises and customer complaints to proactively addressing issues before they become problems, thus improving the experience of your customers. In addition, service sensors keep you apprised of cleaning

activities so you can monitor and enforce your service level agreements with your vendors and staff.

Desk / Meeting Room Occupancy

As a tenant, it is almost impossible for you to know if you have leased the optimal square footage and implemented efficient usage programs to house your operation. You just can't be everywhere all the time to observe space utilization. How well is a hot-desk program running? Where are available desks? Do we really need this many meeting rooms?

Occupancy sensors give you real-time availability information as well as historical usage over time. Real data provides evidence needed to right-size your real estate footprint, potentially saving up to 25% in rental or leasing expense.

Healthy Office

As co-working spaces are disrupting the competitive landscape for office tenants, facilities managers are under increasing pressure to ensure tenant workers are happy, comfortable, and delighted. Workers are demanding amenities and refuse to settle for office space that is sub-par. Healthy environments are not only essential to keep tenants, they are also cost effective.

Real-time monitoring of temperature, humidity and CO_2 levels leads to real energy savings across a floor, a whole building or multiple buildings. Consolidating data over multiple facilities over time provides a comprehensive view of trends and performance, including the impact of external forces such as weather that leads to more informed and better decisions.

IoT Solutions for LoRaWAN

The value of the LoRaWAN network will be realized as new solutions are adopted and deployed within expanding IoT ecosystems. Leaders from businesses in every industry in every geography can make better decisions with insight derived from data. You don't have to be an IoT expert to take advantage of these digital transformation initiatives. You just have to get started.

Figure 4 – Microshare Dashboard

About the Author

Charles Paumelle is a Microshare Co-Founder and leads both our global marketing effort and our international business development. An evangelist for the Internet of Things, Charles was chosen as the Co-chair Marketing for the LoRa Alliance, the fastest-growing Low-Power Wide Area Networks (LPWAN) organization. Prior to helping launch Point.io and then Microshare.io, Charles led the Accenture practice for business process management (BPM) across Europe, Africa, and Latin America. He was also Managing Director International for Knowledge Rules, a global BPM consultancy. Charles held a number of eBusiness/eCommerce leadership positions for Dell Computers, and Shell Gas LPG. Charles started his career whilst still at university working for Pegasystems (NASDAQ: PEGA) where he helped grow their business in new sectors and geographies.

About the Company

Microshare is a patent-pending Data Sharing Platform that delivers disruptive operational savings and insights across all industries leveraging IoT and Data. The Microshare Platform drives Digital Transformation and Data Monetisation including new revenue streams, new business models, and strategic differentiation by delivering a unique data governance fabric to the Internet of Things revolution.

https://www.microshare.io/

6.6 Accelerating IoT Solutions Deployment in Healthcare with myDevices' End-to-End Platform

By Jennifer Noble, myDevices

LoRaWAN technology enables myDevices to accelerate the creation and deployment of complete end-to-end IoT solutions that automate refrigeration monitoring and more in healthcare facilities.

Introduction

Healthcare providers are responsible for ensuring patient safety at all times, and proper refrigeration plays a vital role in the process. Medications, vaccinations, laboratory specimens, and cafeteria foods are all temperature sensitive, and prolonged temperature variations can produce harmful or even fatal bacterial growth that can spoil thousands of dollars of inventory within hours. The Joint Commission (TJC), Food & Drug Administration (FDA), Centers for Disease Control (CDC), and other regulatory agencies require healthcare providers to monitor refrigeration temperatures regularly and maintain detailed records proving that temperature-sensitive inventory has been stored properly. Inspectors can demand to see those records at any time. Requirements are not standardized across regulating agencies, however, making accurate record-keeping a challenge for all healthcare providers. If records are inaccurate, or inspecting agents otherwise find that temperature records violate compliance requirements, hefty fines can be imposed. IoT solutions developed using the LoRaWAN standard such as myDevices' IoT in a Box quickly, efficiently, and reliably automate refrigeration monitoring processes, enabling healthcare providers to ensure proper refrigeration and regulatory compliance.

Solution Architecture

Many healthcare providers monitor refrigeration units manually, but temperatures can fall out of range at any time, and it can be difficult to accurately determine whether or not the inventory is safe enough to keep. LoRaWAN wireless temperature sensors and gateways are combined with the LoRaWAN network and application layer to protect inventory around the clock with real-time, automated remote monitoring. The sensors first collect refrigeration temperature data and transmit it to the gateway. Next, the gateway sends the data to the network, and finally the network sends the data to the IoT in a Box application server for evaluation. If temperatures fall out of the predetermined range, the application server will send an alert notification to designated contacts via text message, so they can take immediate corrective action and safeguard temperature-sensitive inventory.

Cost Optimization

When it comes to deploying IoT solutions in healthcare facilities, particularly in hospitals, cost optimization is a big concern. Facing increasing pressure to reduce expenditures while

improving operational performance, administrators are looking to LoRa-enabled solutions to ensure low-cost efficiency with low-power, long-range capabilities.

LoRaWAN devices are wireless, which helps to minimize upfront hardware and installation costs. There are no wires or cables to run, no holes to drill, and solutions can be installed within minutes by simply placing sensors and gateways in the areas to be monitored. Low power consumption helps to minimize long-term solution costs. Unlike Wi-Fi and cellular devices, LoRaWAN devices can hibernate in between data readings to conserve power and extend device battery life to 20 years or more.

Coverage is also an important factor in minimizing solution costs, as refrigeration units often span multiple departments and floors. Without adequate coverage, sensor data may not transmit properly to the network. Wi-Fi-based solutions typically provide only a few hundred feet of range indoors, which means larger facilities will have to purchase multiple gateways and install them in various locations to ensure each sensor transmits data properly. And while cellular-based solutions offer greater range than Wi-Fi, their high bandwidth needs leave them best suited to outdoor applications. The LoRaWAN specification has long-range capabilities. A single LoRaWAN gateway enables IoT in a Box to provide up to 15km (9.3 miles) of deep coverage for both indoor and outdoor applications.

Similarly, trying to get a signal out of a refrigerator or freezer that has heavy insulation can be difficult. Wi-Fi-based sensors cannot be placed directly into refrigeration units, for example, because doing so will block the network signal. Instead, Wi-Fi based sensors must be tethered to less than optimal locations, such as the outside edge of refrigeration units, and a wired probe fed into the door seal. Repeated door openings and closings can cause the probe to deteriorate over time. Since LoRaWAN sensors are wireless, they can be placed directly into refrigeration units without interrupting network communication.

Customization Beyond Refrigeration

One of the biggest benefits of the LoRaWAN standard is its plug-and-play interoperability among applications and IoT solution providers. Devices and gateways communicate bi-directionally over the LoRaWAN network. A single gateway can relay millions of messages from class A (lowest power), class B (deterministic downlink latency), and class C (lowest latency) end devices, enabling the creation of any use case. In other words, end devices can be mixed, matched, and scaled to meet unique and dynamic IoT needs.

Such plug-and-play versatility is particularly useful for healthcare facilities, where monitoring needs vary greatly and often extend beyond refrigeration. Some facilities are large and want to deploy temperature monitoring sensors across hundreds of refrigeration units, while some are small and need only a few sensors. Some facilities want to monitor pipe systems for water leaks, while others want to monitor air quality, motion, parking, or occupancy. The possibilities are endless.

How Data is Secured

At the edge, data is secured using LoRaWAN Network Security using AES-128 encryption keys as described in IEEE 802.15.4/2006 Annex B. MAC for encryption. These keys are used for layer related frames encryption using a Network Session Key (NwkSKey), while the applications frames are encrypted using an Application Session Key (AppSKey). As with routers, the gateways used in a LoRaWAN network can forward from and to the cloud using Datagram Transport Layer Security (DTLS), Transport Layer Security (TLS), or a Virtual Private Network (VPN) without decrypting or knowing any of the keys.

Within the LoRaWAN Network Server (LNS) lies the registry of devices and gateways that transmits data to the cloud. Each device is identified with two unique identifiers called the DevEUI and AppEUI. The AppSKey and NwkSKey are negotiated along with the regional regulation settings when the device joins the network. Then, the LNS forwards all traffic to the myDevices Cloud using HTTPS (TLS/SSL).

The myDevices Cloud is an IoT platform that supports different device connectivity protocols without compromising security. All communications, including database connections and service-to-service communications within the platform, are encrypted. Also, data is encrypted within the database, and we enforce a rigorous access control list (ACL) for staff, personnel, and customer data. The customer applications and backend system utilize the standard OpenID Connect (OIDC) and Security Assertion Markup Language (SAML) 2.0 protocols to initiate and authenticate any requests to our APIs. These protocols are the modern security standard for user-to-service and service-to-service communications.

LoRaWAN Technology in Action

The Director of Nutrition at Saint Luke’s Hospital in Kansas City, MO reached out to inquire about automating temperature monitoring in cafeteria refrigerators in order to ensure regulatory record-keeping compliance and to protect high-risk patients from contracting potentially fatal foodborne illnesses. The Director had staffed two full-time nurses at a cost of more than $19,000 per month solely for the purpose of manually recording temperatures in over 100 refrigeration units, six times per day. Each manual reading took between one and two minutes to record, however, which means the nurses dedicated a combined total of no fewer than 300 hours in an average month recording more than 18,000 temperatures. The nurses unsurprisingly struggled to keep up with the required readings, jeopardizing patient safety and placing the hospital at risk for costly regulatory violations. We deployed a total of three LoRaWAN gateways and 112 wireless LoRaWAN sensors in the span of one afternoon to instantly and automatically record temperature and humidity data around-the-clock.

LoRaWAN technology enables sensor data to be transmitted in real-time to the IoT in a Box web and mobile applications, enabling the nurses to remotely monitor all refrigeration units at any time of day or night. This also means that instead of dedicating each working hour to recording refrigeration temperatures, the nurses can get back to more meaningful tasks such as tending to patients. Notifications were set to instantly alert designated contacts via SMS text message if temperatures fall out of the predetermined range, so they can immediately correct issues before they become costly problems.

LoRaWAN technology enables our refrigeration monitoring solution to deliver substantial time and money savings while increasing monitoring protection, beginning from day one of deployment. LoRa-enabled sensors capture nearly 200,000 more readings every month at a cost of less than $1,000 per month, placing the hospital on track to save more than $223,000 per year just in monitoring costs.

About the Author

Jennifer Noble is the Content Creator at myDevices, where she creates a multitude of meaningful written and visual marketing communications to drive brand and product awareness across a variety of channels. Prior to myDevices, she worked as a marketing associate for McMaster-Carr Supply Co, an industrial supplies distributor. Her professional interests include tech startups, entrepreneurship, creative marketing, and social media. She received her Bachelor's degree in English Literature from Rutgers University and her Master's degree in Communication Management from the University of Southern California.

About myDevices

myDevices, the IoT solutions company, empowers system integrators, MSPs, VARs, and enterprises to quickly customize, deploy, and commercialize LoRaWAN-enabled IoT solutions. myDevices is the creator of IoT in a Box™– finished and customizable IoT solutions for a variety of vertical markets; Cayenne – the industry's de facto IoT Solution Builder; and IoT Ready Program™– trusted catalog of drag-and-drop IoT devices. Our mission is to simplify the connected world by providing tools that accelerate the deployment of device and connectivity agnostic IoT Solutions for commercial refrigeration, smart buildings, smart agriculture, asset tracking, and other IoT verticals. myDevices is headquartered in Los Angeles, CA.

https://mydevices.com/

6.7Building end-to-end network with IP over LoRaWAN

SCHC, the new open standard for compression-fragmentation, opens new horizons for IoT and accelerates LoRaWAN adoption. It allows the use of Internet technologies in the LPWA networks and offers neutrality, rapid service development, and deployment, plus independence from the underlying radio technology.

By
Matthieu Brient, Acklio
Philippe Cola, Bouygues Telecom
Ana Minaburo, Acklio
Pascal Thubert, CISCO

LPWAN limitations and pains

The Low-Power Wide-Area Networks (LPWANs), such as LoRaWAN, NB-IoT, SigFox, and LTE-M are a set of long-range communication technologies dedicated to IoT deployment.

A low-power, wide-area network creates a point-to-point connection between a device and an application server. Many LPWAN applications and implementations are tightly coupled to the LPWAN technology underneath. Thus, reaching for new markets and integrating additional technologies means laborious re-engineering of both cloud and the device applications.

These emerging network technologies do not follow the Internet model based on IP protocols. They are not interoperable with each other and they are challenging to integrate with existing architectures and information systems. Each IoT project involves first to choose a technology, then to adjust the choice of devices and platforms and to manage integration with the existing architecture. Thus the market remains somehow organized into technological silos.

To harness the full power of LoRaWAN and other LPWAN technologies, it is critical to couple them with other ecosystems. It will guarantee smooth inter-working and interoperability, and enable common components for management and security, as well as shared application profiles.

Lessons learned from the Internet and the hourglass model

In the Internet world, interoperability and multi-connectivity are native properties. Indeed, the Internet Protocol is open, allowing every player to operate around a universal standard reference. The Internet uses a layered model. Each layer implements a protocol and abstracts the complexity of the other layers. This hourglass model isolates the application from the network which saves significant engineering efforts for every actor in the value chain. In its core, IP provides optimized services such as addressing and routing, services that any application can inherit without reinventing the wheel.

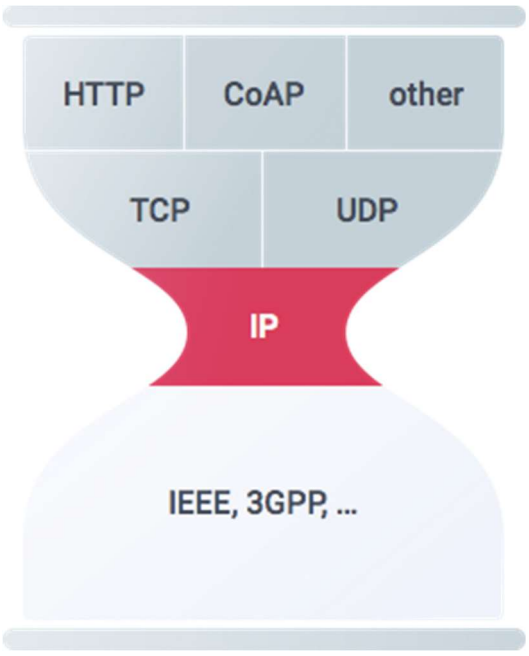

Figure 1: Internet convergence modeled as an hourglass

Due to the severe bandwidth restrictions, until recently, LPWANs were not able to carry the greedy IP stack of the classical Internet, or of less-constrained IOT networks. The newly developed innovation and the resulting standards - object of this article - is discussed later. Plus, unlike the Internet, the Internet of Things lacks governance around a global model. The risk is to see technological actors work in isolation by adapting technologies to suit particular use cases. Instead, we believe in a standard convergence solution to alleviate the complexity that comes with this growing diversity. That is an hourglass model at a layer above the radio, like what IP does for the Internet.

Standardization work at IETF

We are committed to standardization at the IETF[1] with this very objective of providing the IoT community with standard and sustainable solutions. The IoT-oriented IETF working groups already produced the first wave of mature standards for compression and fragmentation in the early 2000s: RoHC (Robust Header Compression) in 2001 and 6LoWPAN (IPv6 over Low-Power Wireless Personal Area Networks) in 2007.

RoHC, a header compression mechanism

[1] By defining the open standards that guarantee the interoperability and evolution of networks, the IETF is the standardization organization behind the success of Internet. The mission of the IETF is to make the Internet work better by producing high quality, relevant technical documents that influence the way people design, use, and manage the Internet. www.ietf.org

The RoHC header compression scheme was designed for point-to-point multimedia channels to reduce the header overhead of multimedia flows using IPv4 or IPv6 or IPv4/v6/UDP headers. It diminishes the size of the transmitted header by removing redundancy in the header packet and the flow of information. It relies on a shared context, which does not require any state, but the packets are not routable. The RoHC compression starts by classifying header fields according to their changing pattern. This classification is used to generate the context, then shared between the two end-points.

RoHC header compression does not work for the constrained nodes of LPWAN networks. However it may provide the necessary level of compression for LPWANs, the compression depends on a steep learning phase that depends on the network traffic and the data flows. Plus, the context synchronization and update processes do not take into account the energy limitations and the transmission rate of LoRaWAN.

6LoWPAN, a fragmentation mechanism

6LoWPAN was defined to optimize the transmission of IPv6 packets over LLNs (low-power and loss networks) as IEEE802.15.4. Here, the Edge router manages the data exchange between devices inside the 6LoWPAN network and communications to the Internet.

6LoWPAN reduces IPv6 and UDP headers overhead by eliding header fields when they can be derived from the link layer. It assumes that some of the header fields will frequently carry expected values. 6LoWPAN provides stateless header compression that can omit part or all of the IP addresses. But the level of compression that can be achieved depends on whether the IP addresses of source and destination match the MAC addresses. RFC 6282 also enables a stateful address compression based on context, which can be seen as a step towards SCHC; in that case, and for selected UDP ports, the IPv6+UDP headers may be compressed down to 5 octets.

Because IEEE 802.15.4 operates its own reliability mechanism by retransmission (ARQ), 6LoWPAN does not have reliable delivery. Some LPWAN technologies do not provide such acknowledgments at level 2 and would require other reliability mechanisms.

Work within the LPWAN working group and SCHC

In response to the lack of a solution applicable to the specific constraints of LPWANs, a new group has been created within the IETF en 2016. The “LPWAN Working Group” is working on the management and operation of constrained-node networks, security and lifecycle management, and semantic interoperability. It focuses on describing a common architecture between the different LPWAN technologies — for example, the very sporadic traffic, the low throughput, and the need to optimize battery life.

LPWAN technologies generally share a very similar structure: a radio-layer gateway connects the end-device, and a network-layer gateway aggregates multiple radio gateway and enables

connectivity to the applications. They also share the common desire to enable IP technology between the network application and the end-device application, to offer more portable services and more generic tools.

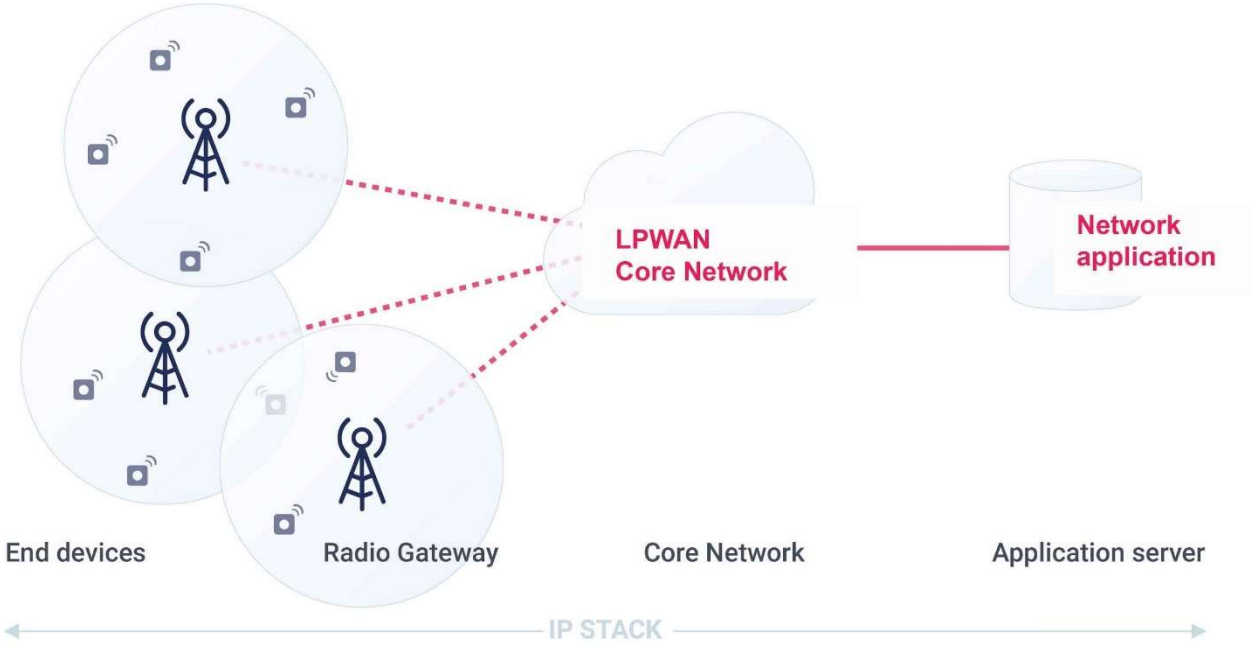

Figure 2: Low power, wide-area network common architecture

A large part of the work done at the IETF LPWAN working group was to identify common needs for functionalities in the LPWAN gateways to standardize the protocols. This generic architecture makes it possible to consider the adaptation of the Internet protocol regarding any LPWAN technologies.

To provide a compelling solution, it aimed to provide the most efficient IP compression and fragmentation to address any LPWAN technologies, like LoRaWAN. IPv6 has been designed to allocate addresses to all endpoints connected to the Internet. Nevertheless, the header overhead of at least 40 bytes introduced by the protocol is incompatible with LPWAN constraints. Based on the experience acquired through the development of RoHC and 6LoWPAN, they propose a new framework, SCHC, that addresses these constrained networks and devices.

Header compression and fragmentation tailored to LPWANs

Static Context Header Compression (SCHC – pronounced “chic”) is the new IETF framework that solves the header overhead problem by reducing the header size. It combines the advantages of RoHC context, which offers a significant level of flexibility in the fields processing, and 6LoWPAN behavior to elide fields that are known from the other side.

How does “Static Context Header Compression” (SCHC) work?

The SCHC compression and fragmentation is specifically tailored for LPWANs. The compression is based on a static context shared and stored in both the device and the network sides.

This new framework takes advantage of the characteristics of these constrained networks (no routing, known traffic format) and reduces the impact of protocol headers to a few bits. With SCHC, the devices and the network keep a set of rules to describe the communication context. The compression uses these rules for compression and decompression. Since the content of packets is highly predictable in LPWAN networks, the context rules can be provisioned beforehand, reducing the need for more network traffic.

The “static context” indicates that the description of the rules does not change during the transmission. It avoids complex resynchronization mechanisms and the learning phase that is incompatible with LPWAN characteristics. In most cases, IPv6/UDP headers are reduced to a small identifier. The context is stored in the end system of the host. The network gateway can learn through a provisioning protocol during the identification phase as the encryption key.

Each context possesses a list of rules, which describes the header fields. Each field is described with a value, an operator (corresponding or not to the rule), and an action (how to compress/decompress the field). A rule only describes the behavior for the header fields and does not give the packet format that is already known on both sides.

The rule has an identifier, which will be used to determine which rule is used. If a rule corresponds to a header’s values, it is used to compress and the values are not necessarily sent in the link. Since the contexts are synchronized, reading the rule’s value is sufficient to reconstruct the value of the field at the other end. In other cases, the value of a field must be sent in the link.

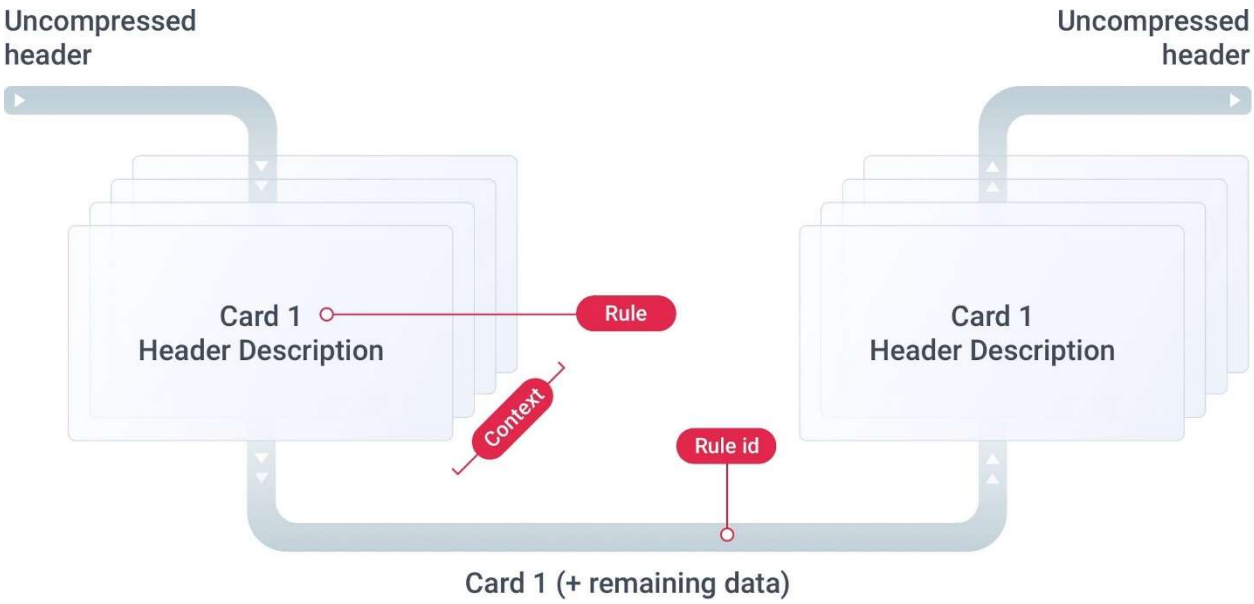

Figure 3: Static Context Header Compression global view

CoAP compression using SCHC

CoAP (Constrained Application Protocol) is an implementation of the REST architecture for constrained devices. It is specifically designed to use minimal resources, both on the device and the network. The combination of Internet Protocol and CoAP releases the full power of the IoT. On a security level, just as HTTP is secured using Transport Layer Security (TLS) over TCP, CoAP is secured using Datagram TLS (DTLS) over UDP.

CoAP supports the naming mechanism, the resource management and the manipulation of primitives between a client and a server. The strength of CoAP resides in its full compatibility with HTTP. It is also possible for a CoAP application to interact with an HTTP server through generic gateways and vice versa. This allows objects to be anchored in a dominant ecosystem built around REST principles.

Nevertheless, the size of CoAP header was incompatible with LPWAN. To compress and decompress CoAP header, SCHC uses the same framework as for IPv6/UDP. It simply adds some compression and decompression actions, that are more adapted to the CoAP behavior.

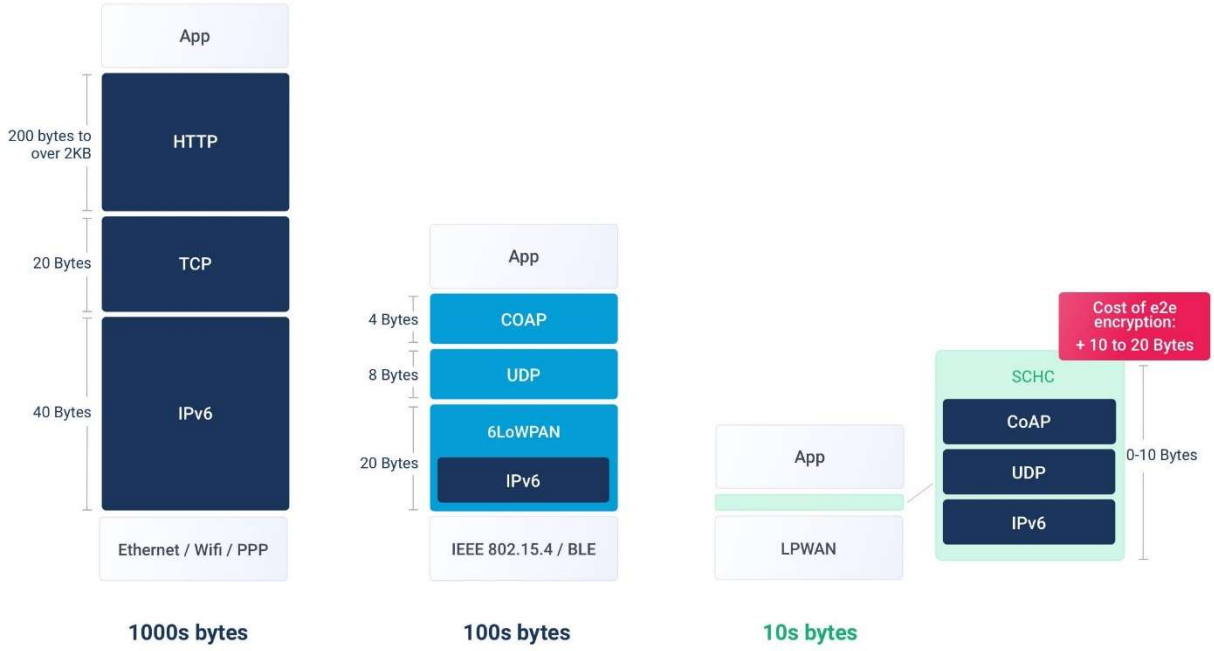

Figure 4: Internet Protocol stack compared to 6LowPan and SCHC compression frameworks on IPv6/UDP/CoAP

Security

IoT devices, and especially LPWAN-connected devices will have a major role in the operation of the big infrastructures of tomorrow. As such, any information coming from them may be used to influence the decisions of how these infrastructures are operated. In addition, controlling the

behavior of these infrastructures will also be an integral part of the use of these LPWAN connected devices. In that sense, IoT devices must be protected and all communications - secured - with the state-of-art security mechanisms.

The stakes are nothing less than having a city infrastructure become unresponsive, the fully-automated machines on the factory floor irreversibly damaged (as with the example of Stuxnet), location and presence data leaked to hostile entities, etc. among the multitude of threats that endanger all aspects of IoT networks and services that will be imbedded in the infrastructures surrounding our daily lives. LPWAN devices and networks may be particularly vulnerable to these kinds of threats because of their very strengths - long battery life makes way for devices that remain untouched for many many years, potentially being integrated deep in the surroundings, and long range communication allow for attacks over big distances, lowering the cost of large-scale attacks.
It is therefore crucial that communication between the IoT device and its application be available at all times, whether the application is in its initial commissioning, regular operation or decommissioning phase. Also, unwanted communication from/to the IoT device should be barred at all times to limit the effects of a compromised device to the outside world and limit the capability to compromise more devices remotely.
The Internet is the most hostile environment in the digital world. A computer connected directly to the Internet becomes a target of scanning and attacks in less than several minutes, and without protection gets infected almost instantly. Yet, we do our online banking over the Internet, we perform our daily busyness, use secure file-sharing, and video-conferencing. How is that possible? The most hostile environment is the place where only the most robust security mechanisms can survive. The strength and reliability of these mechanisms is the reason we can use the Internet for all these daily uses. Any other technological choice is doomed to face tough challenges ahead - as the true test of it will come years down the road, when it becomes popular, and when it will be too late to change.

One robust and straightforward way to provide such protection, enabled by SCHC, is to isolate the device and its application in an overlay network. Furthermore, hostile Internet-originated threats can be mitigated by operating on Unique Local Addresses that cannot be reached from the outside of the overlay. Also, security controls exist for network communications at each layer of this IP stack. Dedicated IoT standards, such as DTLS or OSCORE, ensure end-to-end encryption even for constrained networks and devices. Tested and verified time and time again, they provide an answer to the question - will my network and my devices be secure in two years' time? In five? In ten? Will my services be vulnerable to quantum attacks? Will my customers be obliged to replace all devices deployed over a large area, embedded in Cyber-Physical Systems, because hackers around the world are making daily ransom demands to unblock their operation?

IoT without security is dangerous IoT. Security cannot be improvised and cannot be added as an after-thought. The only way to provide secure system operation that withstands the test of time is to use proven end-to-end mechanisms tested and in operation in some of the most

challenging environments today. Internet-stack based security brings all this, and much more, and is now a viable option thanks to SCHC.

New IoT use cases made available by SCHC

SCHC as a universal adaptation layer

SCHC drastically reduces LPWAN's constraints and enables to build rich and energy-efficient services. Every developer knows how to develop with IP. Then SCHC implementation avoids the custom developments constrained by the LPWAN niche technologies and accelerates the deployment of new IoT devices and services on a large scale.

Moreover, the use of LPWANs are now within reach of many more players throughout the IoT value chain. Actually, thanks to SCHC-based protocol adaptation, most IoT use cases can be operated on LPWANs. SCHC is for example a strong candidate for DLMS and Wireless MBus over LoRaWAN for smart metering, LWM2M over LoRaWAN for device management, or other proprietary protocols.
LoRaWAN to connect native IP equipment

Also, now that we are able to communicate on IP over LPWANs thanks to SCHC, we can implement LoRaWAN as redundancy link for out-of-band management or supervision link.

LoRaWAN Backup Connectivity Kit for Cisco routers

Implemented for the first time by Acklio, SCHC enables interoperable use of the IP protocol suite across any LPWAN technologies. On February 2018, Acklio and Cisco implemented IPv6 operation over LoRaWAN. Together, we demonstrated how SCHC enables secured SSH communications over LoRaWAN at the LoRa Alliance All Members Meeting.

The scenario showcased addresses the worst-case outage by offering a redundancy solution for the network administrator. It allows them to connect, diagnose and fix problems remotely. For example, in the case where the main communication link fails between SCADA industrial controller and a remote utility grid (Ethernet or 3GPP), LoRaWAN could be used as a backup link. Here, SCHC allowed to set up an additional virtual wire to a Cisco IOx network devices.

To connect the IP equipment via LoRaWAN, a dongle ensures the compression/decompression of IP messages to send them over the radio link. These messages are interpreted by the Acklio LPWAN IP Core before being transmitted to your third-party applications. Full end-to-end SSH communication is ensured between the IP equipment and the application.

LoRaWAN-connected after-sales service

After the first successful worldwide experience between EDF, Acklio and Bouygues Telecom / Objenious explained at the LoRa Alliance’s “All Members Meeting” in Vancouver (June 2018),

Bouygues Telecom continued the work with its partners to set up the LoRa Back-up connectivity. For an operator point of view, this allows greater responsiveness.

Now, Bouygues Telecom connected after-sales service allows remote technical diagnosis in the event of a customer breakdown, or WAN link is down. Here, the implementation of SCHC enhances the redundant administration solution by using the Objenious LoRaWAN network. It allows operation teams to open a remote administration console via SSH for the first level of troubleshooting, even in the event of a failure of traditional operating networks.

Conclusion

Leveraging existing LPWAN technologies, with SCHC, the IETF value proposition is aligning the very different radio technologies on a common hourglass model. This model involves an extremely compressed form of IPv6 and CoAP between the end-device and the network gateway. It provides common management of the gateways and enables secure, Internet-based services to the applications.
Using IP as a common language removes the need for bespoke bridges to translate network traffic from one technology to another. SCHC removes the risks of vendor lock-in and provides future-proof investments. It addresses specifically the IoT ecosystem needs and accelerates the IoT market as a converging LPWAN Architecture.

But most of all, IP is a mature open standard. Ubiquitous, most (if not all) operating systems support it and it provides native interoperability with current information services and network infrastructures. Simple, efficient and scalable, IP is ready to connect billions of devices.
SCHC over LoRaWAN provides a native solution with no compromise on security or scalability.

About the Companies

Acklio is a French startup created in 2016. The cofounders impulsed a technology that ensures interoperability and interconnection between any type of networks and radio technologies. This innovation, called SCHC, is recognized as an international standard by the IETF. Deploy it today with Acklio's software, the first market-ready implementation.

Bouygues Telecom (and Objenious is subsidiary) have deployed the first the LoRaWAN Network in the world in 2016 with national French coverage (4500 LoRaWAN Gateway).

https://www.ackl.io/

6.8 How multi-mode location on LoRaWAN helps create business value

By Brad Bush is Managing Director, Longview IoT

LoRaWAN is known for enabling low-power sensors to send data for many different applications over long distances. However many believe that the processing which accurate location measurement demands creates too much battery drain, despite the efficiencies and informed decision-making which location ubiquity promises. This article will show how inside and outside tracking technologies of various types (BLE, Wi-Fi SSID, LoRa RSSI, TDOA, and GPS) have various battery, accuracy and feature trade-offs.

By Brad Bush, Managing Director, Longview IoT

In the future, location ubiquity will be the norm. We will know where everyone and everything is at all times. This will not seem strange or different, it will just be part of life, as computing ubiquity became the norm in the 1980's with PCs, connectivity ubiquity became the norm in the 1990's with the internet, and mobile ubiquity became the norm in the 2000's with mobile phone technology. We are already on our way to location ubiquity with mobile phones, cellular, and in-vehicle tracking advancing on one side, while Bluetooth and radio frequency ID (RFID) and other product/warehouse-based location technologies advance on the other.

With all these technologies, a power vs distance trade-off is present. Cars and mobile phones are mostly powered when using the global positioning system (GPS) or are quickly recharged if power is consumed. Have you ever wondered why your battery drains faster when using map applications? A standard phone GPS satellite lock takes between 12-30 seconds and a full GPS satellite update could take as long as 12 minutes, (https://www.theverge.com/2018/8/17/17630872/smartphone-battery-gps-location-services) a real problem when tracking small things. We need devices that can still be tracked but use much less battery power. Figure 1.

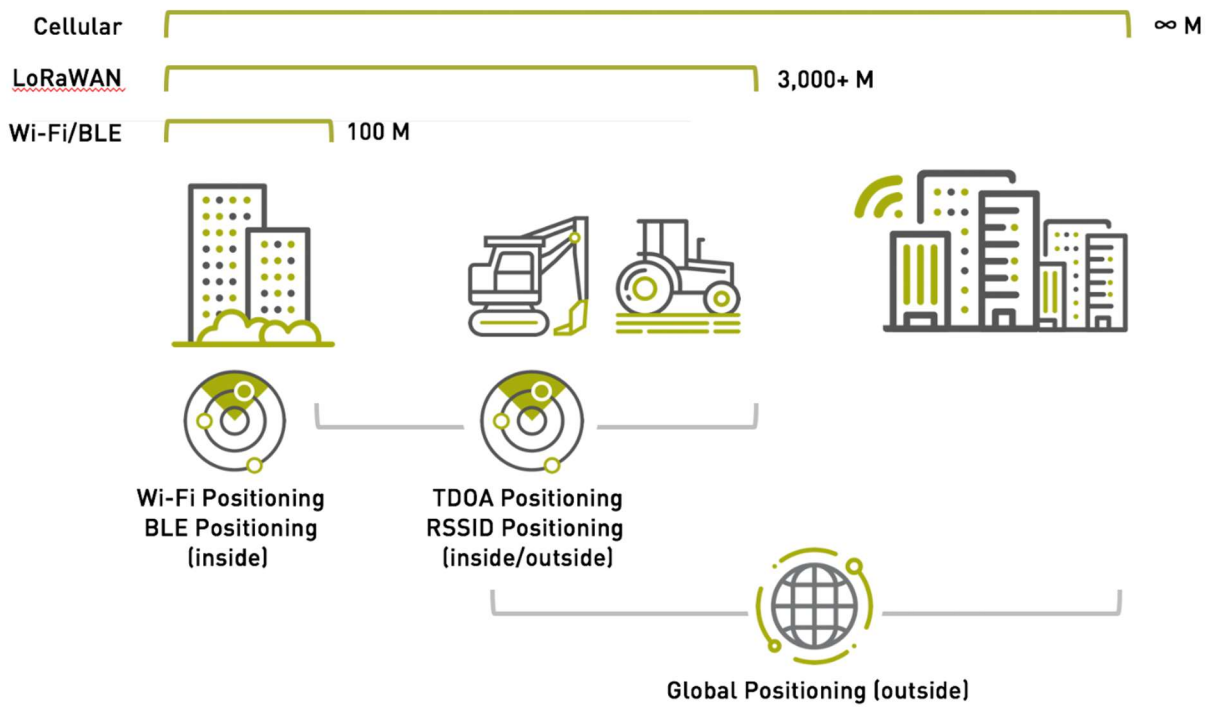

Figure 1: There is a trade-off between power and distance in network technologies.

When location measurement happens within close range, as is the case in a warehouse, there are other factors in play. RFID has dominated this use case due to its ease of use. RFID requires a reader which emits a magnetic field and detects “tags” that are within the field area – the reader could be a door or a scanner or some other device. RFID tags come in two flavors, passive without a battery and active with a battery. Active RFID tags can go up to 30m from the reader but cost more than passive tags. The issue with this technology is that it is proximity related. If you don’t have a reader in range, you don’t know where an item is. It is not an answer to our location ubiquity problem, just a piece of the puzzle.

So what low-powered technologies are available that can cover ranges greater than 30m? There are several factors to consider when looking at location technologies. First, what is the accuracy that you need – is it within 3m or is 30m good enough? Second, how often do you need to measure the location – is it near real-time or once a day for example? Third, how will you determine and send the location using the least amount of power?

Let us start with the third question first. LoRaWAN is a low powered wide area network (LPWAN) technology that has gained great popularity due to its low cost, high performance, and support/development of its open standard by many organizations via the LoRa Alliance. LoRaWAN allows devices to send short messages at low power across long distances (up to 10km but practically around 3km). This is very useful in filling in the location ubiquity gap. LoRaWAN can be paired with several technologies that can measure different geolocation accuracies. Some companies are even combining several location methods to get better results.

Let's start from the most power-heavy and accurate technologies and work our way through to the least-power heavy and least accurate. As stated above, GPS is the most common service used in measuring outside location. GPS is a system of around 30 satellites, run by the U.S. government. A device can locate which satellites are in the sky and where they are and use this information to determine its position. Because you have to be able to see the satellites, GPS only works when outside with a clear line of sight to the sky. GPS accuracy can range from 30cm to 10m.

Several companies such as Actility and Longview IoT have introduced GPS-based LoRaWAN devices. The key issue in using GPS is battery consumption, as a standard location fix takes 1-3 minutes. LoRaWAN can easily be used today on GPS-based devices for tracking once or twice a day on common batteries that will last for years. The above companies have worked on proprietary algorithms to reduce the time needed to locate the device. These technologies make it possible both to track in near real time, such as every five minutes, and to have battery life of up to ? months], a combination that makes applications like tracking people and equipment outside very viable.

Time difference of arrival (TDOA) geolocation is another common method for locating assets. This method uses the time it takes for a signal to move from the device to several gateways (base stations). It calculates the difference of the times it takes the message to get from the device to the varying gateways and using triangulation math, determines the location of the device. There are a couple of key factors that are required for TDOA to be most accurate. First, the entire system must be in very tight time sync (all the gateways and end devices). Second, the geometry of a TDOA system works best in a tight triangle structure. The more linear the gateway alignment, the less accurate the TDOA readings will be. Increasing and densifying the number of gateways will increase accuracy, but only to the limit of the system timing accuracy. TDOA systems can be setup using timing from GPS chips or other chips. A LoRaWAN TDOA system is able to get within sub 50m accuracy at a slight additional power consumption.

LoRaWAN has its own built in location method based on received signal strength indicator (RSSI). RSSI is the amount of power received by a gateway from a given sensor endpoint (typically measured in dBm). Power signals decay in a non-linear way from the base station. This means that when you are far away from a gateway, you will get much less accurate readings. RSSI works better with more gateways and at closer spacings. There has been much written about the accuracy of RSSI measurements, but LoRaWAN RSSI is able to get around 100m accuracy with dense gateway configuration. One of the major advantages of RSSI is that it is just a measurement of the strength of the signal, therefore, nothing else new is happening and no additional battery is being used.

Both TDOA and RSSI can be used for indoor location, but the results vary greatly. With machine learning algorithms and training, a building could be mapped by "learning" the signal strengths or time delays of a specific space, but as people and furniture move around it would affect these measurements. There are many companies such as Everynet that use these technologies in

basic geolocation services on inexpensive devices with long-life batteries – the tradeoff is accuracy.

To augment location services in indoor location Wi-Fi and Bluetooth are being used in conjunction with LoRaWAN and the above methods. Wi-Fi is probably the most common way to "assist" location services today. When you turn on your phone's assist mode, it looks at all the nearby available Wi-Fi networks and their various signal strengths via the Service Set Identifier (SSID). This information is fed into a global database of all the Wi-Fi access points and their exact locations (it is really a form of RSSI done in the cloud). Using this information your phone is then able to get a relative location to assist in GPS and use less power. This is great when you are in a big city with tall buildings blocking the satellites and lots of Wi-Fi access points around.

LoRaWAN systems are now also using Wi-Fi location. They only have to have a Wi-Fi chip that looks for the networks and measures them, they don't even have to get on a Wi-Fi network. The surrounding Wi-Fi networks are then passed back via LoRaWAN to the cloud to combine with GPS or other data. This can fill in gaps when a device is inside or tighten up accuracy of TDOA or RSSI. The great thing about Wi-Fi service when it is used for location is that it does not take very much battery power. Several companies such as Longview are creating multi-mode location services with GPS and Wi-Fi.

Bluetooth (also known as Bluetooth Low Energy or BLE) is another technology that has been combined with LoRaWAN to increase indoor accuracy. Bluetooth beacons are small devices that can be placed around an indoor area such as a warehouse or a retail shop. Then when Bluetooth-enabled devices are near, they can "listen" to the messages from the BLE beacons. Usually several beacons are placed in a room and it uses a form of RSSI to determine position very accurately within the room. Beacons or devices can be enabled to communicate their location back over LoRaWAN. This can reduce network connectivity setup and operational costs inside a large building like a warehouse where normally wireless or hardwired networks would be required to reach each beacon.

Different use cases have different accuracy and power requirements. If you are monitoring a person for fall detection, you may need to check them every few minutes at minimum. If you want to see where a hospital bed is, it may be acceptable to check only once every day. If you want to know where you left your jack hammer on a jobsite, you probably need sub 10m accuracy, but if you want to know where your cows are on your 1,000-acre ranch, 100m accuracy may be fine. These tradeoffs and use cases mean that all of the above methods for location services are viable.

In many cases, it makes sense to combine GPS, TDOA, RSSI and Wi-Fi location into a multi-mode location service. This would mean that a single device could be configured to use the best service for the use case and act intelligently as required. Outside and able to see satellites, the device would use GPS. It could fall back on TDOA, RSSI or Wi-Fi when needed under various

conditions. These multi-mode location services are going to be more popular as they increase while using the least amount of battery.

Location based services will be used in many use cases, from tracking assets such as tools and equipment, to monitoring worker safety. We will also see location services tied to other sensor types such as temperature and location which are currently being used to ensure full cold-chain monitoring so that produce is kept cold throughout its life cycle.
LoRaWAN, combined with many current geolocation methods, can put us firmly on the path to true location ubiquity by filling in the gaps between cellular- and local-based services. This will create lasting business value, as everything is more valuable when you know where it is.

About the Author

Brad Bush is Managing Director, Longview IoT at the communications innovation company Carnegie Technologies, and a Senior Partner of Fortium Partners. Brad has more than 25 years of technical and leadership experience, serving in roles as CIO, COO and CMO for high-growth, technology-based companies. Brad understands the human side of technology, using his technical and business background to execute near-impossible scenarios and create new business models. He is a thought leader in many technologies and was named a WebRTC pioneer, an InformationWeek CIO top innovator and a top 50 CMO in telecom. Brad has held various roles at Carnegie including Chief Strategy Officer and Chief of Staff. Before joining Carnegie Technologies, Bush was Chief Operating Officer at Dialexa, a research, design and creation firm which builds innovative software and hardware platforms including Internet of things, drones, artificial intelligence and virtual reality. Leading Dialexa's sales, marketing and operations, Bush worked on technology projects for Fortune 500 companies and high-growth startups.

Prior to joining Dialexa, Bush held several positions at telecom software maker GENBAND, including six years as Chief Information Officer, two years as Chief Marketing Officer, Chief Integration Officer, head of sales operations, and Chief-of-Staff reporting to the CEO. Before GENBAND, he was responsible for business unit operations at Tekelec Switching Solutions Group, which was acquired by GENBAND in 2007. Prior to Tekelec, he spent eight years at Nortel Networks, where he held various strategic positions in information technology and wireless operations. An entrepreneur at heart, Bush founded Sharpsite Interactive, a successful startup specializing in Internet application development.

He holds a Bachelor of Science in Civil Engineering from Rice University and an MBA from Southern Methodist University. When he is not working, Bush enjoys playing music and traveling with his family.

About the Company

Longview, a Carnegie Technologies Company, is for companies that want an IoT solution that just works, without the complexities and struggles of getting sensors, networks, software, and hardware from different vendors to work together. The enterprise-grade system can scale for use across multiple sites with hundreds of thousands of sensors and was built on the open standard LPWAN LoRaWAN™ protocol to cover vast terrain. Longview's triple layer security builds both hardware-based keys and an additional software certificate layer over the already robust LoRaWAN security, while its Super-B protocol enhances the standard LoRaWAN protocols by allowing for more efficient bandwidth usage and field over-the-air (OTA) upgrades. Large-scale asset tracking is required across many industries today, and Longview is the most secure, reliable and efficient solution to solve this difficult problem. For more information or to schedule a demo, visit www.LongviewIoT.com.

Carnegie Technologies was founded in 2010 to develop mobility products and services that connect people, things and networks throughout the world. The Company has a deep technology heritage and delivers groundbreaking connectivity and convergence solutions for mobile operators, satellite communications, Internet of Things, corporate board governance, mobile entertainment and much more to fundamentally transform the way we think about communications, public and personal safety entertainment and economic efficiencies. Headquartered in Austin, Texas, the Company has offices and product development on four continents.

https://www.carnegietechnologies.com/

6.9 Use of LoRaWAN Gateways Multi-Network SIMs in the Airtime Tariff Applications

Multi-network (aka roaming) Subscriber Identity Modules (SIMs) have often been proposed as the solution to intermittent or unreliable Global System for Mobile Communications (GSM), communications, but this is only half the story – a story rarely understood or fully explained by those selling such products

By David Evans, Robustel

Introduction

In this article, we explore the technical and commercial considerations of using multi-network SIMs to make backhaul from cellular-based long range (LoRa) gateways more reliable than when a single network SIM is used (Figure 1).

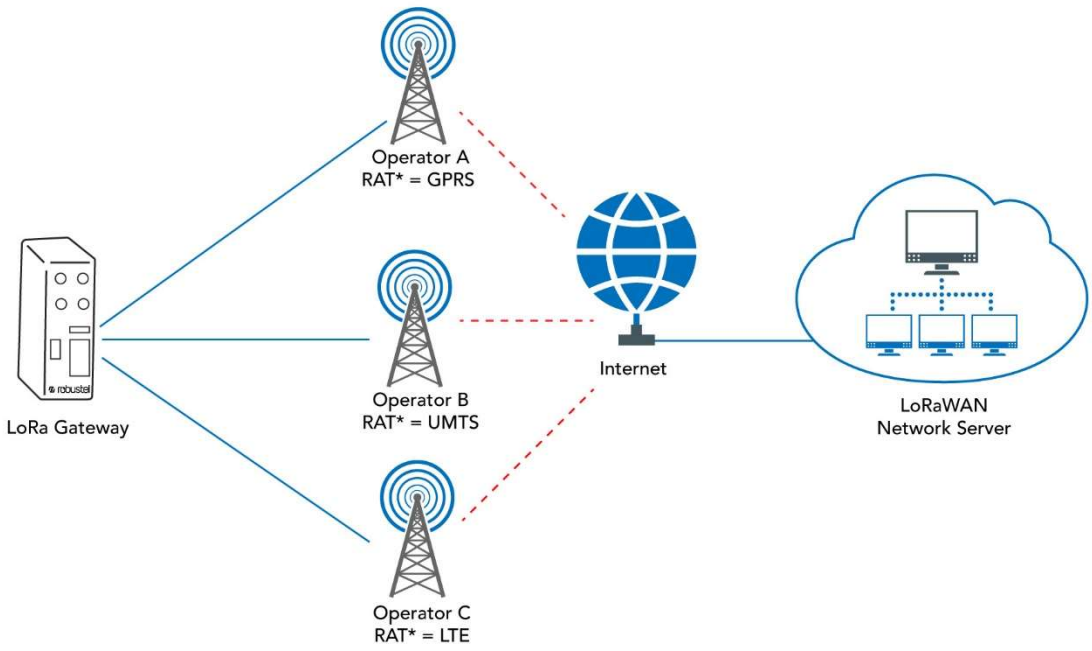

Figure 1: Using a roaming SIM provides diverse routes to the internet giving a higher chance of successful communications between Gateway and Server

SIM types

To fully understand the options, we open with an overview of fundamental SIM types that can be used in a long-range Wide Area Network (LoRaWAN) gateway.

1. **Single Network**

   Single network SIMs cover a single network, typically in a single country. If no coverage exists on a particular network at an installation location, especially if it is in a building, then the communications will not work. And if the deployment is multinational, then it becomes necessary to maintain lots of different SIM agreements for different territories, which is logistically difficult and virtually impossible to manage at scale.

   Roaming or multi-network SIMs, such as those from aggregators KPN, Vodafone GDSP, Tele2, Telenor, and their value-added reseller partners, have been the traditional solution proposed to the problems just mentioned. What most people don't realize is that not all roaming SIM cards are equal.

2. **Steered Roaming SIM**

   This type of SIM can in theory access multiple networks per country, but in practice will favor a specific network or subset of networks for the mobile operator's commercial benefit. Usually steering is not implemented for the benefit of reliable communications but to make the scheme less costly by using only preferred networks. But this use of preferred networks results in the opposite of what IoT service/system providers who need as close to 100% network uptime as possible want, because the favored network may be unable to deliver a service when a non-favored one can!

   While steered roaming is a complex subject just touched on here, one solution is to select a good quality unsteered roaming SIM provider to remove the vagaries of steering from the application.

3. **Unsteered Roaming SIM**

   As the downsides of steered roaming have become more visible in the market, certain mobile network providers have started to provide and champion unsteered roaming SIMs. This means the SIM has no preferred network list onboard and does not suffer operator steering at the network level. Essentially, an unsteered roaming SIM is a blank canvas that simply provides access to networks as and when required with no 'loading of the dice.'

Having established a source of a good quality unsteered roaming SIMs, we can then investigate how a LoRa gateway reacts to such a SIM and other measures to optimize gateway-to-server communications reliability.

Network selection process

One of the most misunderstood concepts surrounding roaming SIMs is that of network selection.

Contrary to what many people have been told, SIMs do not play an active part in network selection (USIM toolkit applications excepted). Choosing an appropriate network is a function of the hardware, not the SIM card.

LoRa gateways are built using off-the-shelf cellular (typically 4G LTE) modules such as those manufactured by Telit, Gemalto, Quectel, Huawei, and other companies.

The module takes care of all interaction with the mobile network from the RF layer upwards. Most modules have a TCP/IP stack onboard for simplified application development but can also act in a modem mode, whereby the TCP/IP stack in the host system (typically Linux for a LoRa gateway) acts as the IP endpoint for gateway/server communications.

The host system communicates with a cellular module using industry standard attention (AT) commands, as defined in GSM07.07 and GSM07.05, as well as a range of (module) manufacturer-specific AT commands.

These commands determine how the module in a LoRa Gateway should behave, and a subset of AT commands are directly associated with the network selection process.

The vast majority of LoRa gateway manufacturers will use automatic network selection by issuing “AT+COPS=0” to the module, which instructs the module to take care of registering to a network per the set of rules in TS 3GPP 23.122 as detailed below:

“The MS (cellular device i.e. LoRa Gateway) selects and attempts registration on other PLMN/access technology combinations, if available and allowable, in the following order:

i. *either the HPLMN (if the EHPLMN list is not present or is empty) or the highest priority EHPLMN that is available (if the EHPLMN list is present);*

ii. *each PLMN/access technology combination in the ‘User Controlled PLMN Selector with Access Technology’ data file in the SIM (in priority order);*

iii. *each PLMN/access technology combination in the ‘Operator Controlled PLMN Selector with Access Technology’ data file in the SIM (in priority order);*

iv. *other PLMN/access technology combinations with received high quality signal in random order;*

v. other PLMN/access technology combinations in order of decreasing signal quality."

To aid with understanding of the automatic network selection process, please see the Glossary of Terms in the table below.

PLMN = Public Land Mobile Network = Mobile network
HPLMN = Home Public Land Mobile Network = Home network
EHPLMN = Effective Public Land Mobile Network = Network with effective status of Home network
Access Technology = Radio Access Technology such as GPRS/UMTS/LTE approximating to 2G/3G/4G

It's important to note that the automatic network selection process does not make mention of the "last known good network" paradigm, whereby knowledge of previous successful registration to a network is stored in the SIM. This paradigm will take precedence and means a SIM will attach to the previous network, irrespective of a change in location or change in network conditions, so long as basic criteria are met.

When using an unsteered SIM without a notification of a previously registered network, none of conditions *i* to *iii* above are met, meaning that the module will take action *iv* as the first step in automatic network selection.

Action *v* entails the module performing a network scan to identify local networks and associated technology (GPRS/UMTS etc.).

Once complete, the module will ascertain which networks have a good quality signal and choose *at random* the network where it will attempt registration first, selecting from among all good quality networks.

The key term here is *at random*.

Many people peddle the 'strongest signal' myth, suggesting that a roaming SIM will always go to the 'strongest signal' first. As can be seen from the TS 3GPP 23.122 specification discussed earlier, registration on the strongest signal will only be attempted in step *v* if there is no possibility to register to a network using the methodology described in step *iv*.

If one imagines 200 people disembarking from an airplane at a holiday destination and switching their phones on at roughly the same time, it is quite evident why a random and distributed sharing of the roaming traffic is desirable, hence the fundamental design of the automatic network selection process. (The roaming process was originally invented for consumer handsets, not LoRa gateways!)

The problem with automatic network selection

Now that the process of automatic network selection is understood, it's easy to discover its limitation. Automatic selection only uses signal strength to determine a usable network, but a low-level received signal strength indicator (RSSI) measurement is by no means definitive proof that the chosen network can support end-to-end IP based communications. So, we need something else, something that will supplement automatic network selection to make sure we are getting the best service available.

Commercial considerations of roaming / multi-network SIMs

Having ascertained that a roaming SIM will likely give a better "uptime" for an estate of LoRaWAN gateways, it is essential to consider the financial implications of a move from single network to roaming SIMs.

Until recently, roaming SIMs have only been commercially viable for very low data usage applications such as vehicle tracking and smart metering—especially when considering an unsteered roaming SIM, where any network can be freely visited without limitation.

However, changes in the market now mean that on some continents, roaming is now very affordable, and the EU is certainly leading the charge on this subject.

$10 to $20 per GB per month is now available in the EU for a good quality unsteered roaming service. Is it worth paying extra to get closer to 100% uptime across an estate of LoRa gateways? While the answer depends on the intended application, it's important to know that those whose backhaul is depending entirely on cellular communications have a choice to receive better service and what they must do to capitalize on this choice.

It is this fundamental issue that Robustel's SMART Roaming capability aims to solve, using the simplicity of automatic network selection but supplementing it with health-checks and the use of manual network selection to make sure reliability is as high as it can possibly be.

SMART Roaming – a solution to an age-old problem

Robustel's LoRaWAN gateways support SMART Roaming, which can be configured to check for loss of mobile data communications on the current network and force a change to an alternative network within a short period of time. This can save the cost of site visits and provide peace of mind that the highest possible reliability cellular methodology is keeping the packet forwarder of the gateway in communication with the LoRaWAN network server.

SMART Roaming checks not only signal strength, but also ping times and ping completion to build a more complete picture of the current connection. If the health-check is failed, the router will dynamically assess the quality of alternative networks and change to the next best if communications are lost or are of a 'low quality.'

This fundamental concept is used to great effect in the metering market, where millions of devices are deployed, and an overall estate uptime of close to 100% is achieved as a consequence. This technology is generally developed and tuned for a specific cellular (2G/3G/4G) module and roaming SIM combination. Generalizing the solution for a multitude of networks is not easy, but that's what has been achieved in Robustel's LoRaWAN gateways courtesy of the SMART Roaming application.

Summary

The idea of being able to transmit data on multiple radio networks rather than depending on just one as a means of improving reliability has long been recognized. However, it is a complex subject matter that falls between the domain of the SIM provider and the hardware provider. Breaking down these knowledge barriers and providing an off-the-shelf SMART Roaming solution is Robustel's contribution to helping with the proliferation of LoRaWAN as a standards-based and highly flexible LPWAN solution.

About the Author

David Evans has 20 years of experience in technical sales and marketing with a focus in industrial and wireless automation. Much of this experience has been gained working on projects in the IoT / M2M space across a multitude of vertical markets.

A self-confessed geek at heart, David has excelled within his chosen roles due to a fundamental desire to understand the complex and find business benefits & opportunities where others don't care to look.
An education in Physics as well as a comprehensive education in life at the University of Sussex in the late 90s provided David with the tools to understand technology and be able to present to and share with others in a meaningful way.

About the Company

Robustel are one of the world's leading manufacturers of industrial quality solutions for the IoT and M2M market. Robustel's portfolio of award-winning solutions are comprised of: Wireless Modems, Routers, Gateways, EDGE Computing, Cloud Software and End-to-End IoT solutions. Robustel are committed to helping businesses and industries across the world continue to solve their IoT & M2M problems with robust, secure, scalable and creative solutions from hardware to complete 'IoT in a box' services. We are Robustel, and our goal is to Make Things Connected.

https://www.robustel.com/

6.10 Device-to-Cloud LoRaWAN IoT Solution

Edge Computing with LoRaWAN provides significant benefits for multiple industries. Whether it be high speed process control systems or autonomous feedback loops, an Edge Computing system provides enhance performance, reliability and security. The following article highlights key capabilities and benefits of LoRaWAN Edge Computing applications in manufacturing, oil & gas, transportation, and lighting.

By Nik Kitson, CEO, Haxiot

What is Edge Computing?

Edge Computing is the ability to perform functions with partial or complete autonomy from the Cloud. Edge Computing is typically invoked in situations where both sensor and control functions are utilized autonomously using a distributed software agent. LoRaWAN was designed as a cloud-based centralized technology without support for Edge Computing. Companies such as Haxiot have developed multi-site Edge Computing that enables LoRaWAN to be deployed within industries that have different requirements than traditional cloud services can offer.

Why is Edge Computing useful?

Different use cases create a variety of benefits. The simplest edge computing benefit to understand is bandwidth. Many sensors generate vast quantities of data, while distributed edge computing agents look for exceptions. Video or sensors with extremely high sample rates that generate terabytes of raw data each day can be analyzed in real-time to output digital data. For example, HD video cameras on highways can send number plate data messages rather than streaming gigabytes of data to a centralized cloud system.

Latency

High speed process control systems, such as those found in manufacturing or Oil & Gas require extremely fast response times to report events in high performance systems. Several failures are detected and require a response in millisecond time frames to avoid damage to equipment. High latency or internet disruptions rarely occur, but unavoidable events in any public cloud system. For example, Haxiot X-ON edge system has deployed edge-to-edge latency using LoRaWAN sensors as low as 100 milliseconds, compared to 1000’s of milliseconds using cloud.

Remote Sites

Either the cost or availability of cloud connectivity limits the impact and effectiveness of cloud-based solutions. Sporadic cloud connectivity exists in transportation systems, satellite coverage, or underground / underwater systems. Deployments with these characteristics require many of the cloud functions to run autonomously using Edge Computing. Using Haxiot X-ON cloud with MODBUS intelligent endpoints connected over LoRaWAN an industrial customer was able to

poll 10 samples per second, forwarding alarms and compressed polling data to the cloud. The polling granularity was improved by 6000x and while reducing bandwidth consumption by 96% over traditional MODBUS monitoring.

Feedback Loops

Sense & Control fully-autonomous feedback loops are increasingly being used to operate even more complex systems. Early systems were based on triggered thresholds with fixed responses, which used static rules based on experience. Modern Sense & Control systems are moving to cloud-generated machine learning algorithms, which deploy on edge computing nodes that automatically establish non-linear triggers using deep learning patterns from large cloud data sets. The algorithms apply these hyper-local sensor datasets to dynamically control local systems. These systems have a wide range of responses beyond simple thresholds, as they have trained on extremely large data sets.

Use Cases

Manufacturing

Process Control systems in manufacturing are designed for reliable, granular control of real-time mechanical systems like motors, pumps, production lines, and robotic assembly tools. These systems are often influenced by factors outside of the system sensor suite and need external systems to monitor external variables, making operational decisions on the fly. Vibration sensors in motors are often retrofitted on or adjacent to high speed spinning motors utilizing advanced sensors that provide real-time alerts when diminishing performance is detected. Edge Computing software can take intelligent vibration data and make real-time adjustments to the engine speed to extend until the maintenance window or to immediately shut off and move to a redundant system.

Using LoRaWAN enabled wireless factory sensors, production lines and be instrumented with a range of sensors. It is critical that this data is secured and inter-connected at the edge where the wired process control and wireless sensor systems intersect.

Oil & Gas

The unique nature of the Oil & Gas industry is its large number of highly diverse assets, which are distributed over many geographical sites. This creates challenges of accessibility, connectivity, and data integrity. The vast majority of Oil & Gas assets have not been historically instrumented with Internet of Things data connection devices. The data generated by these devices has several key values to the ROI of Oil & Gas assets:

1. Real-time analytics on the edge of the network for process control and predictive maintenance. Safety and Security monitoring of critical infrastructure
2. Long-term trend analysis to optimize operations

Oil & Gas operations generally have slow and/or expensive access to cloud services due to the nature of where the infrastructure operates. Cloud-based services are less reliable and responsive to the needs of real-time responses and Safety/Security considerations.

LoRaWAN sensors can instrument remote assets cost effectively, providing real-time data simultaneously to both edge applications as well as cloud analytics.

Transportation

Data collection in motion is important for moving goods around the globe. The growing importance of product-level sensors that have very low power and memory footprints requires edge computing platforms to collect and store IoT sensor data while goods are crossing the "digital deserts" where connectivity is unavailable. Since the majority of the oceans and land lack cellular or wi-fi coverage, the ability for edge computing systems to authenticate, secure and cache sensor data allows more detail on products in transit, rather than receiving product information only at origins and destinations. The transportation industry uses this method extensively to optimize truck and train maintenance activities using location data for all assets combined with temperature data to more accurately predict failures.

Critical to this industry are technologies that can be deployed with distributed intelligence moving with the assets and sensors. Providing LoRaWAN sensors with edge compute enabled gateways allows data services to operate with intermittent cloud-access. This provides customer benefit for creating time-complete audit trails for asset maintenance or compliance.

Lighting

Commercial outdoor lighting is undergoing a transformation from simplistic on/off with trigger-based light sensors to complex light, color, and predictive maintenance. The next evolution of commercial outdoor lighting controls is coming to market as a fully-integrated lighting and data business, generating insights into the environment with built-in environment sensors, such as pollution and CO monitoring. Traditionally, cloud-based solutions involve significant latency issues for sensor and local switch-controlled lighting systems that are commonplace today. Non-data driven analog sensor integration combined with directly-connected control systems have provided a simple but inflexible means of hyper-local control. Edge Computing with wireless Sensor & Control systems have better response time and reliability for lighting control rather than using cloud-based systems.

The evolution of LoRaWAN has closely followed the evolution of cloud technologies. It is notable how little traction these cloud-based approaches have achieved in displacing "Edge Compute" systems, which are dominated by legacy technologies, such as MODBUS or Zigbee. A revolution is coming to edge computing for those who are able to adapt to the new breed of centrally-deployed LoRaWAN solutions with high scale, distributed edge solutions that better meet the needs of industrial and commercial IoT customers.

About the Author

Nik Kitson, is a seasoned technology expert across Internet, telecommunications and Internet of Things for over 20 years. Started his career in the service provider industry culminating with senior technology and strategy in roles UUnet and Vodafone. Transitioning to Silicon Valley, Nik lead mobility strategy and architecture for Juniper Networks. Lead business development and strategy for Cisco's largest account AT&T account. Rolled out the first city-wide LoRaWAN network in Atlanta for the Cisco-AT&T Foundry innovation lab before co-founding Haxiot. Developed first distributed Edge Computing solution for LoRaWAN.

About the Company

Haxiot provides LoRaWAN solutions for enterprise & industrial. The Haxiot X-ON hybrid cloud/edge platform provides rich data value-added services on edge gateways and cloud, supporting real-time process control & predictive maintenance. Haxiot's IoT solution is proven to enable enterprise customers with faster time to market, lower integration cost & a positive ROI. Haxiot is headquartered in Dallas, TX with distributors in China, North America & LATAM markets.

www.haxiot.com

6.11 The Modernization of IoT Networks

By Dave Kjenda, CTO, Senet

In recent years, we've seen a combination of technologies converge to become the foundation of change across the network landscape. Cloud computing has become more dynamic in its capabilities, creating new opportunities for application and service creation and delivery. Virtualization, microservices and adaptive security models have also become key elements of modern network design as the digitization and connectedness of our world moves forward at a rapid pace.

These cornerstone technologies have also been driving factors behind the growth of the Internet of Things (IoT). As we move toward a world in which everything is connected, it is no surprise that the legacy networks designed primarily for high-bandwidth applications have proven to be sub-optimal in price and performance for many of the emerging machine-to-machine (M2M) applications that need to send and receive small amounts of data while consuming almost no power. The changing architecture of these networks has also meant a change in the way they need to be managed. With dynamic performance management and application optimization becoming increasingly important, network management is being re-invented and is motivating a range of companies to offer new protocols, pricing structures and customer engagement models suited to low-bandwidth M2M Internet of Things (IoT) communication.

Next Generation OSS/BSS Platforms — Built for IoT

Key components of network connectivity and management, often transparent to end users, are the Operational Support System (OSS) and Business Support System (BSS) platforms. Focused on the network and services, an OSS is typically used by network planners, service designers, and engineering teams and orchestrates and automates the 'back-office' network management functions. Business Support Systems comprise a separate set of applications supporting commercial, revenue, and customer-relationship activities. Combined, OSS and BSS deliver the full set of capabilities a service provider needs to operate a network and sell services.

Being able to securely connect, activate, and monitor IoT devices at massive scale, in a multi-tenant and multi-vendor environment, across a broad range of applications, is the new standard for network operators. Network providers need to be able to manage the OSS/BSS features of the network server, packet core, data streaming, security, performance of the Radio Access Network (RAN) and End Device adaptive data rates (RF tuning). In addition, the IoT application management environment provided by the network operator must efficiently enable gateway deployment and provide scalable, secure, end-device onboarding, application service provisioning, and visualization tools.

While some operators struggle to transform their monolithic legacy systems to support these requirements, many are gaining a competitive advantage by opting to have their network managed by a network services provider with an OSS/BSS built for IoT. At a time when IoT applications are being envisioned and built daily, what should operators consider?

Open Standards-Based Technology: The OSS/ BSS must combine innovative and open standards-based technologies, enabling the IoT application provider to leverage the operational experience of the network provider. This partnership between the network provider and IoT application provider mitigates much of the risk often associated with the adoption of new technologies.

Built for the Scale of IoT: The OSS/BSS offering must be secure, cost effective, and scalable, providing operational efficiencies that scale to support billions of connected devices. It must also support technology that enables the highest levels of network reliability and service level excellence. Backing by a responsive support organization ensures the IoT provider's service offering is predictable and reliable wherever it is being offered.

Built for Market Expansion: The capabilities and services should make it possible for IoT application providers to rapidly and economically move from concept to pilot to massive-scale commercial deployments. The network offering must be capable of providing rapid 'time to coverage', meeting the application provider's goals for expansion beyond their traditional areas of service.

As-a-Service Cost Efficiencies: An OSS/BSS designed with the optimal set of functionalities for operating a LPWA network will always have a significant cost advantages derived through a scalable cloud deployed "as a service" product offering.

Organic IoT Network Expansion

Just as carriers are moving away from legacy OSS/BSS technology in order to support the scale and operational complexities of IoT, they are also recognizing the need for new engagement models across the IoT ecosystem to gain a competitive advantage. The "if you build it, they will come" network deployment and customer engagement models designed for personal mobile communications are not well suited for scaled machine-to-machine communication. Replacing them are new models designed to grow through partnerships and application expansion. As a result, new business opportunities for companies of all types to participate in the IoT economy are emerging.

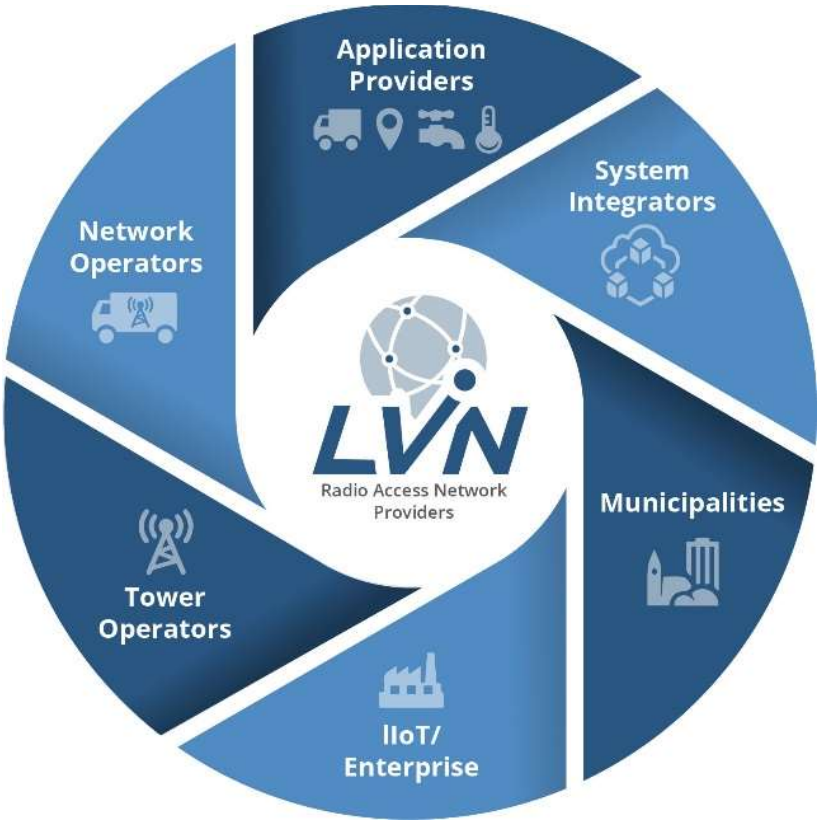

Figure 1: LVN Radio Access Network

With this model, application and solution providers can connect to a network where it is available and then partner with the network operator to contribute to the buildout of the network by purchasing and deploying low cost LPWAN gateways in areas where additional coverage is required. Similarly, system integrators can strategically build their IoT business by providing their growing base of customers with low cost connectivity services when and where it is needed on an application by application basis. In each of these examples, the network is managed by the operator and those contributing to the network build out can benefit from a revenue share based on role they play in the larger network ecosystem.
This same model extends to network operators who have yet to establish an IoT practice or have determined a combination of IoT connectivity services is the best approach for them. By sourcing network management services from a proven IoT partner, cellular, cable, wireless and fiber providers can augment their connectivity portfolio by deploying LPWAN gateways on their existing tower or building assets and extend their branded services well beyond their existing footprint. Figure 1.

Lastly but equally importantly, this model creates opportunities for partnerships between critical infrastructure and service providers - such as municipalities, utilities, cable and cellular network operators and large enterprise organizations - and the citizens they serve. One of the most important opportunities the Internet of Things offers is delivering economic, environmental, and

societal improvements. Whether it is reducing energy consumption in cities, reducing water loss through smart metering and water management solutions, monitoring crop conditions to improve yield or optimizing the cold chain for medical product safety, all parties now have an opportunity to participate in the IoT economy for the greater good of the environment and the welfare of the world's population.

Conclusion

The promise of billions of connected IoT devices is not only coming from expected areas of opportunity such as asset tracking, energy and utilities, and smart cities, but from applications that instrument the ordinary or hidden business activities but yield revolutionary results. Successfully commercializing these opportunities requires transformative enabling technologies and service models including best-in-class OSS and BSS platforms.

Similarly, traditional engagements where service creation and delivery have been a one-way relationship between the communication service provider and the customer no longer apply. Today's businesses seek value-exchange in order to realize the full potential the Internet of Things. This value and success in IoT will be driven by co-operation between network operators, end device manufacturers, solution providers and others looking to proactively and organically expand their roles in the IoT services economy.

About the Author

Chief Technology Officer & Vice President Engineering Dave Kjendal is a senior executive with over 18 years of experience in high-tech innovation management and complex product development. He is a successful leader of large multi-site teams that have delivered dozens of leading-edge LAN/WLAN product lines, embedded and application software designs as well as ODM and OEM development models with responsibility for research and design, product management and product marketing. Prior to joining Senet, Kjendal has held leadership positions with industry-related technology companies – Extreme Networks as Senior Vice President Engineering, Siemens Enterprise Communications as Senior Vice President, General Manager, and as Engineering Architect for Enterasys Networks. Kjendal received his Bachelor of Science in Physics from University of New Hampshire and Master of Science, Electrical Engineering from the University of South Alabama.

About Senet

A Founding and Contributing member of the LoRa Alliance™, Senet is leading the IoT revolution with pioneering experience and expertise in building and operating Low Power, Wide Area Networks (LPWANs). We are working with hundreds of businesses to revolutionize their products and operations by delivering standardized low-cost network connectivity exactly where it's needed, when it's needed and at the right cost.

https://www.senetco.com/

6.12 LoRaWAN: Breaking the IoT walled garden

Just imagine a scenario wherein a Gmail user could send mail only to users who have a Gmail account or one where a mobile phone user registered with Operator 'A' can communicate only with users who are registered with the same operator. Such a walled garden scenario would be a nightmare from a user's perspective, yet this is the case within the Internet of Things (IoT) domain today.

The focus has been on connecting ***things*** to the ***Internet*** using technical features that work only with or between particular apps, devices, servers, etc. This approach creates walled gardens and restricts seamless connectivity in the IoT.

As World Wide Web inventor Tim Berners-Lee has noted, "*I should be able to pick which applications I use for managing my life, I should be able to pick which content I look at, and I should be able to pick which device I use, which company I use for supplying my Internet, and I'd like those to be independent choices."*

By Sandoche Balakrichenan, Afnic

LoRaWAN Stands out from Other IoT Technologies

Unlike many IoT standards, the LoRaWAN standard aligns quite well with the principles of Berners-Lee's quote by enabling users to make their own choices.

For example, LoRaWAN allows for a number of deployment options. There are ***public*** networks having nationwide coverage; ***private*** networks focusing on specific use cases with a limited area of coverage and free ***community*** networks.

As an open standard LoRaWAN allows users to freely adopt, implement, and even extend the LoRaWAN standard for their specific usage. The connectivity over the LoRaWAN radio network is free since it uses a license-free frequency band.

The LoRaWAN ecosystem follows a horizontal model wherein the customers have options in choosing the hardware (end devices, gateway, etc.), the network connectivity (private network, public network, etc.) and the backend (network server, application server, etc.).

In addition, it uses existing protocols that have successfully been used in the Internet, including the Domain Name System (DNS) for features such as Over the Air Activation (OTAA) and roaming. LoRaWAN features differentiate it from other IoT technologies. It closely aligns with the design principles of providing independent choices to the user, thus enabling seamless connectivity, multiple modes of deployment, and continuing evolution to serve tomorrow's business use cases, the connected mousetrap[1], for example.

[1] https://www.xignal.com/products/xignal-mousetrap

LoRaWAN and the Lost Cat

To illustrate how LoRaWAN is different from other IoT technologies, let's consider the real-world problem which ***Cathy***, an amateur technologist and budding entrepreneur, wants to solve using technology. Cathy's cat ***Billy*** gets lost often. A number of times, she has had to put up advertisements in her neighborhood or in the local daily to recover Billy and would like a low-cost IoT service to monitor his movements.

She has heard about Low Power Wide Area Network (LPWAN) technologies and first considers SigFox LPWAN. But her friends and neighbours would also like to monitor the movements of their pet animals. Her entrepreneurial instincts start working. She wants a horizontal platform, which will enable her to provide service to her friends by implementing a private network herself. She is also looking for a DIY type of platform. She selects LoRaWAN because it meets these requirements.

Initially, for testing, she buys an off-the-shelf GPS tracker compact and lightweight enough to be affixed to her pet's collar. The LoRaWAN enabled GPS tracker has the distance – from a few kilometers in dense urban areas up to 15-30 kilometers in rural areas [1]– longer battery strength, and price point (around $50 CAPEX per device) she wants.

Next, she needs a gateway which understands the LoRaWAN protocol and can interpret the data that device, i.e., the LoRaWAN enabled GPS tracker, sends. As she does not want to spend much on her initial set up, she builds a gateway at rather low cost; around $150. At the end, on a budget of around $200, she has the hardware ready to track Billy.

To visualize the GPS data on Billy's movements, she opts to use an open LoRaWAN community network. Her gateway is configured to forward the data received from the GPS tracker over the Internet to the community network backend. To have the complete set-up ready, she completes some basic configuration in the community network backend dashboard. Once the set-up is complete, she is able to track Billy's movements live.

Cathy's neighbors who also want to track their pets' comings and goings need only buy a tracking device that is LoRaWAN enabled and use Cathy's gateway (with her permission) for their connectivity to the Internet. They have to provide her information such as their ***unique*** tracking device identifier and a cryptographic key for authentication. All these data have to be added by Cathy a priori in her account in the community network backend. In turn, Cathy provides her ***unique*** Application ID (provided by the community network backend), which the neighbors then configure in their LoRaWAN device.
Initially, when the neighbor's tracking device tries to make a connection to Cathy's gateway, it needs to send a request asking permission, with the device identifier, the application identifier, and the cryptographic key. As Cathy has already added information about the neighbor's device in her backend, her gateway verifies the credentials and allows the device to be registered in

[1] https://www.postscapes.com/long-range-wireless-iot-protocol-lora/

the backend. Known as Over the Air Activation or OTAA, this is the process of activating the end device before it can communicate on the LoRaWAN network.

The Need for a Globally Distributed Database

By implementing the LoRaWAN set-up as explained earlier, Cathy has established herself as a ***private network operator*** providing LoRaWAN connectivity to her neighbors. For this set-up she has been using the ***backend*** of the community network, which consists of a network server[1], Join Server,[2] and the application server[3].

No special requirement for the OTAA process as long as all Cathy's subscribers plan to use the backend proposed by Cathy. In the event of a subscriber wanting to use his/her own backend, then Cathy's network server must establish a secure session with the subscriber's backend. *Thus there is a need for associating the subscriber's devices to their backend via Cathy's network server for OTAA.*

Similarly, if Cathy or one of her subscribers moves outside of her gateway's geographical coverage, the LoRaWAN enabled devices unaware of the move will be sending messages that could be received by a LoRaWAN enabled gateway in a different network. In order for this gateway to forward the messages received from the device to Cathy's backend, *there is a need for a process used frequently in cellular networks, roaming. The roaming process associates the gateway's network server with Cathy's network server*.

For both associations (either for OTAA or for roaming) discussed earlier, the LoRaWAN standard uses the globally distributed DNS[4] database. DNS helps translating LoRaWAN identifiers, such as NetID and JoinEUI, to the URLs of the respective servers (home NS, and JS). Using DNS for these associations enables LoRaWAN to be completely interoperable on the Internet. It introduces more dynamic usage of the LoRaWAN network, wherein hosting a small-scale private network using LoRaWAN is as easy as anyone hosting a website on the Internet, with a limited deployment cost.

Wrapping it up – Did LoRaWAN Succeed in Breaking the IoT Walled Garden?

The walled garden has been popular since Roman times. They called it *hortus conclusus*[5]*.* This term was introduced in technical innovations wherein the service provider restricts the user's freedom to a certain boundary for using the service.

---

[1] The gateway is configured to forward the received messages to a network server. The network server is responsible for decrypting the received messages and decides whether to forward the payload to a specific application server, Join Server or to drop the message.

[2] The Join Server will be used during OTAA for authentication of the device and generation of the session keys.

[3] Application server is the final destination of the message. The received messages could be stored in a database to be visualized in a manner in which the user requires.

[4] https://tools.ietf.org/html/rfc1034

[5] https://en.wikipedia.org/wiki/Hortus_conclusus

With the advent of the Internet (an open network), many existing walled garden services (in telecom, media, etc.) disappeared. Most of the current IoT standards did not follow the design principles of the Internet, but rather tried to re-introduce the walled garden in their networks. History suggests that it can be a more difficult strategy than in the past.

In LoRaWAN, even though there is a ***necessary***[1] restriction at the radio connectivity end, connectivity at the Internet end is open. The use of Internet standards such as DNS for OTAA and roaming services is a major asset for LoRaWAN, giving it the ability to adapt and increase interoperability with full range of LoRaWAN solutions and services providers. For Cathy, it's giving her choices as described in the earlier description of her pet movement tracking service.

About the Author

Sandoche Balakrichenan received his PhD in Computer Science and Networks from the "Université Pierre-Marie Curie", France." He is currently working as an R&D Engineer at Afnic (the French Network Information Center). His research interest includes IoT identification, interoperability and service discovery. He has been a participant/contributor to several standardisation organisation such as IETF, GS1, LoRa alliance, IoT expert group at the European Commission. Other than IoT, he has also contributed to number of collaborative Internet related projects.

About the Company

Afnic is a non-profit association and the incumbent manager of the ***.fr*** TLD, Afnic is a multi-registry operator of the top-level domains corresponding to the national territory of France, the *.fr* TLD and those of the overseas territories) and of several French projects for new Internet top level domains). Afnic carries out its assignments in the public interest by involving all the relevant stakeholders in its decisions (scientists, the public authorities, and representatives of private sector involved in the Internet in France). As the primary operator in France of registry services on the Internet, the goals set by Afnic are to develop a preference for the *.fr* TLD in France, to help strengthen the resilience of the Internet, and to promote its skills among the Internet community at large.

https://www.afnic.fr/

[1] Necessary - because sending data using the LoRaWAN free unlicensed band requires hardware that understands the LoRaWAN protocol both at the sending and receiving end.

6.13 STM32 Microcontrollers & LoRaWAN®

System development for Low-Power Wide Area Networks (LPWAN) is getting easier and easier. To facilitate acceptance and level the playing field, the LoRa Alliance provides LoRaWAN® standards. Then, developers and chip manufacturers provide possible implementations for hardware and software to make it concrete. It is no surprise that STMicroelectronics, as the leading supplier of 32-bit microcontrollers (MCUs), has teams that have been doing this for its STM32 family of microcontrollers.

By Benjamin Guilloud, STMicroelectronics International NV

LPWAN in general and LoRaWAN in particular are useful for a wide range of applications. Among these, some involve Firmware Updates Over the Air (FUOTA). STMicroelectronics has been developing this capability in a collaboration with Actility, the leader in low-power wide area network connectivity management for the Internet of Things.

Together, ST and Actility demonstrated a FUOTA over LoRaWAN™ Proof of Concept at the 2019 Embedded World Conference. The demonstration was fully compliant with the specifications and recommendations from the FUOTA Working Group of the LoRa Alliance Technical Committee. The implementation that ST showcased at Embedded World ran on a real-world Reliable MultiCast (RMC) production server from Actility.

Unlike a unicast protocol server, which talks to one device at a time, or a broadcast system, where the server indiscriminately talks to all devices, RMC can target particular devices for greater transmission efficiency, without packet losses or out-of-order deliveries. This is the “reliable” element of the multicast server that differentiates it from traditional multicast platforms.

In the framework of this demonstration, Actility’s ThingPark RMC server features two critical modules: one for fragmentation and the other for clock sync. Fragmentation breaks the update down into chunks compatible with the traditional LoRaWAN payload (maximum of 255 bytes). It also uses forward error correction and redundant packets to protect against losses or data corruption. Engineers often overlook this aspect, but properly fragmenting an update to send it over the air is as complex and vital a part of the process as is reassembling the fragments on the receiving end. The clock synchronization is for Class C LoRaWAN devices, where the products are constantly listening to the RF spectrum. By offering a clock sync system, Actility enables FUOTA on Class C devices and makes upgrading firmware over-the-air a reality with existing systems.

From the RMC server, the FUOTA update reaches the LoRa end device. Clock Syncing ensures it can be captured. For the receiving end, ST’s engineers developed an update agent that acts as an intermediary between Actility’s embedded libraries, ST’s Secure Boot and Secure Firmware Update (SBSFU) system (Figure 1).

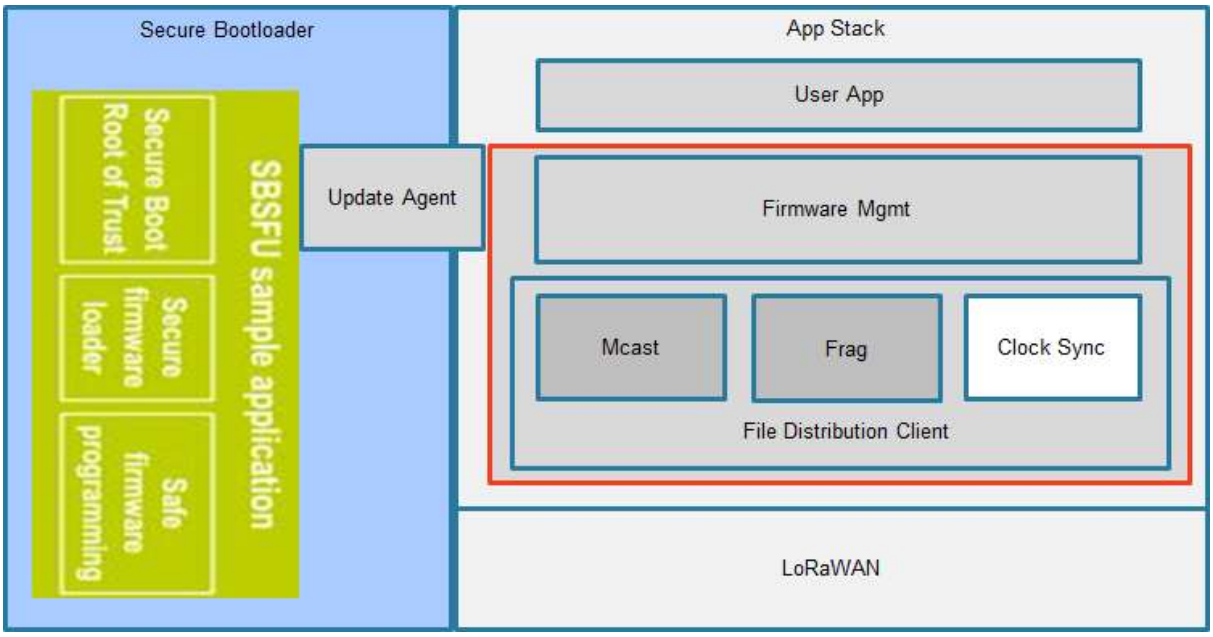

Figure 1: FUOTA library embedded in STM32 MCUs

Once the stack processes the packages and reassembles the fragments from the RMC server, the update agent ensures that the SBSFU system on the microcontroller securely writes the proper information on the right memory space, updating the firmware. Correctly capturing, parsing, and reassembling these small packages on the LoRa terminal is no small feat, and the update agent is a remarkably efficient tool in making this happen without error.

At Embedded World, ST showed a demonstration of the FUOTA using an STM32L4 MCU with only 256 KB of Flash.

A network-server-agnostic version is in development and will be available online soon. Engineers can also use ST's existing LoRaWAN stack to start working on their applications, and ST has posted a tutorial showing how to setup a LoRa node in 10 minutes.

ST's LoRa end devices are built around the company's industry-leading STM32 family of 32-bit microcontrollers. These MCUs are based on the Arm® Cortex®-M core processor. STM32 series MCUs meet the full spectrum of design needs, including high performance, real-time capabilities, digital signal processing, low-power and low-voltage operation, and connectivity. Each microcontroller series maintains full integration and ease of development.

Available free as a software Expansion Package, the LoRaWAN stack for STM32 MCUs has a set of libraries and application examples for ultra-low power microcontrollers which control end devices; the ultra-low power MCUs are ideal for assuring maximum operating lifetimes.

Both the software Expansion Package and the LoRaWAN stack support existing and new radio generation LoRa radio expansion boards from SEMTECH, the developers of LoRa modulation.

The LoRaWAN Expansion Package from ST includes an application that can run on the ST Discovery kit that embeds the LoRa module from Murata. It also supports a USI LoRaWAN technology module embedded in an ST expansion board. This level of support helps jumpstart LoRaWAN development (Figure 2)

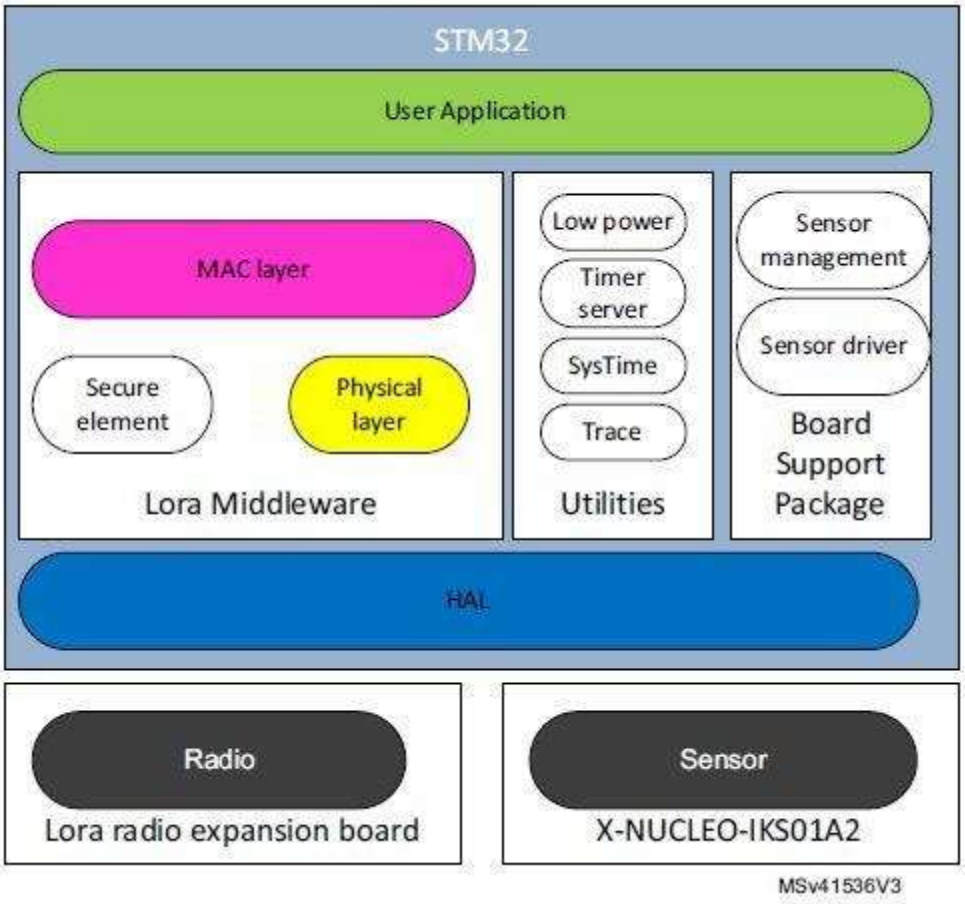

Figure 2: I-CUBE-LRWAN – LoRaWAN library running on STM32 MCUs

LoRaWAN implementation within the ST Expansion Package matches standards requirements such as:

- Compliance with the LoRa Alliance specification protocol LoRaWAN version V1.0.3 March 2018
- Bidirectional end devices with Class A, Class B, and Class C protocol support
- Compliance with LoRaWAN Regional parameters: AS 923, AU 915, CN 470, CN 779, EU 433, EU 868, IN 865, KR 920, RU 864, US 915
- End device activation either through Over the Air Activation (OTAA) or Activation by Personalization (ABP)
- Adaptive data rate support.

Providing complete Hardware and Software Ecosystems along with a LoRaWAN stack implementation ensures that developers can start their development painlessly and with peace of mind while targeting any LPWAN application.

About the Author

Benjamin Guilloud is a Product Line Marketing Manager focused on LPWAN activities targeting applications in the unlicensed Sub-GHz spectrum. He represents STMicroelectronics in the LoRa-Alliance Marketing Committee and in the LoRa-Alliance Roadmap Working Group. Benjamin holds a Master of Science in Micro & Nanotechnologies from Grenoble Institute of Technology. He holds also an International Master's Degree in Nanotechnologies from Politecnico di Torino, Grenoble IT & Ecole Polytechnique Fédérale de Lausanne. He holds also a Master of Science in Micro & Nano-electronics Research from Université Grenoble Alpes.

About the Company

STMicroelectronics is a global semiconductor leader delivering intelligent and energy-efficient products and solutions that power the electronics at the heart of everyday life. ST's products are found everywhere today, and together with our customers, we are enabling smarter driving and smarter factories, cities and homes, along with the next generation of mobile and Internet of Things devices. By getting more from technology to get more from life, ST stands for life.augmented. In 2018, the Company's net revenues were $9.66 billion, serving more than 100,000 customers worldwide.

www.st.com.

6.14 LoRa Technology: Enabling our world to become a Smart Planet

By Semtech

Semtech's LoRa devices and wireless radio frequency technology (LoRa Technology) is a long range, low-power wireless chipset that has become the de facto technology for Internet of Things (IoT) networks worldwide. LoRa Technology enables smart IoT applications that solve some of our biggest challenges: energy management, natural resource reduction, pollution control, infrastructure efficiency, disaster prevention, and more.

Semtech's LoRa Technology has amassed more than 600 known use cases for smart cities, smart homes and buildings, smart agriculture, smart metering, and smart supply chain and logistics. With 87 million devices connected to networks around the globe, LoRa Technology is the DNA of the IoT.

What is LoRa Technology?

LoRa Technology offers compelling features for IoT applications including long range, low power consumption, and secure data transmission. The technology can be utilized by public, private, or hybrid networks and provides greater range than cellular networks. LoRa Technology can easily plug into existing infrastructure and enables low-cost, battery-operated IoT applications. Semtech builds LoRa Technology into its chipsets, which are incorporated into devices manufactured by a large ecosystem of IoT solution providers and connected to networks around the globe. Simply stated, LoRa connects devices (or all things) to the cloud.

Designed for IoT communications, LoRa Technology enables the connection between remote point-of-use devices and low-power, wide area networks (LPWAN) for delivery to analytics applications. LoRa (short for long range) is a spread spectrum modulation technique derived from chirp spread spectrum (CSS) technology. The LoRa physical layer (PHY) is the LoRa modulation itself, as defined by the OSI 7-layer network model in Figure 1. Instead of a cable, the air is used as a medium to transport the LoRa radio waves from an RF transmitter in an IoT device to an RF receiver in a gateway and vice versa.

A homogenous PHY simplifies interoperability and testing at the radio level and offers consistent performance for application developers and gateway manufacturers. It operates in a fixed bandwidth channel (typically 125KHz for uplink channels and 500KHz for downlink channels). LoRa modulation uses orthogonal spreading factors, allowing the network to make adaptive optimizations of individual end-node's power levels and data rates that preserve end-node battery life.

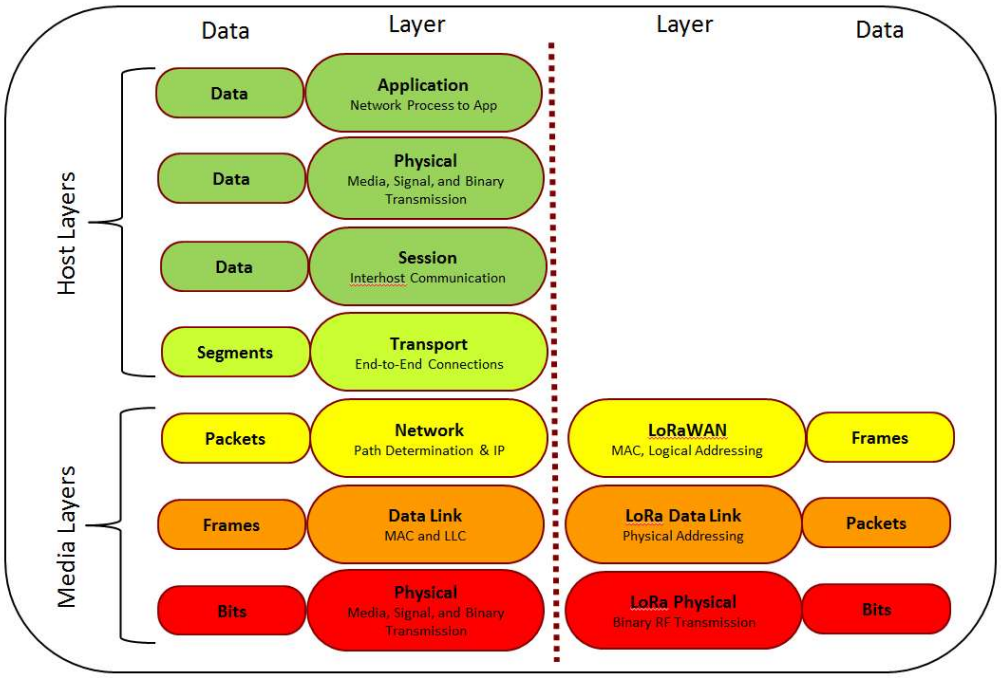

Figure 1: OSI 7-Layer Network Model
© Semtech Corporation. Reprinted with permission

LoRaWAN Protocol

The LoRaWAN open specification is a low-power, wide area networking (LPWAN) protocol based on LoRa Technology. LoRaWAN is an ideal IoT protocol to wirelessly connect battery operated things to the Internet for applications in a variety of vertical markets. The success of LoRa Technology in LPWAN IoT applications speaks for itself: IoT networks based on the LoRaWAN specification have been deployed in more than 100 different countries with an ecosystem supported by more than 500 contributing members of the LoRa Alliance™. While Semtech provides the radio chips featuring LoRa Technology, the LoRa Alliance, a non-profit association, drives the standardization and global harmonization of the LoRaWAN protocol.

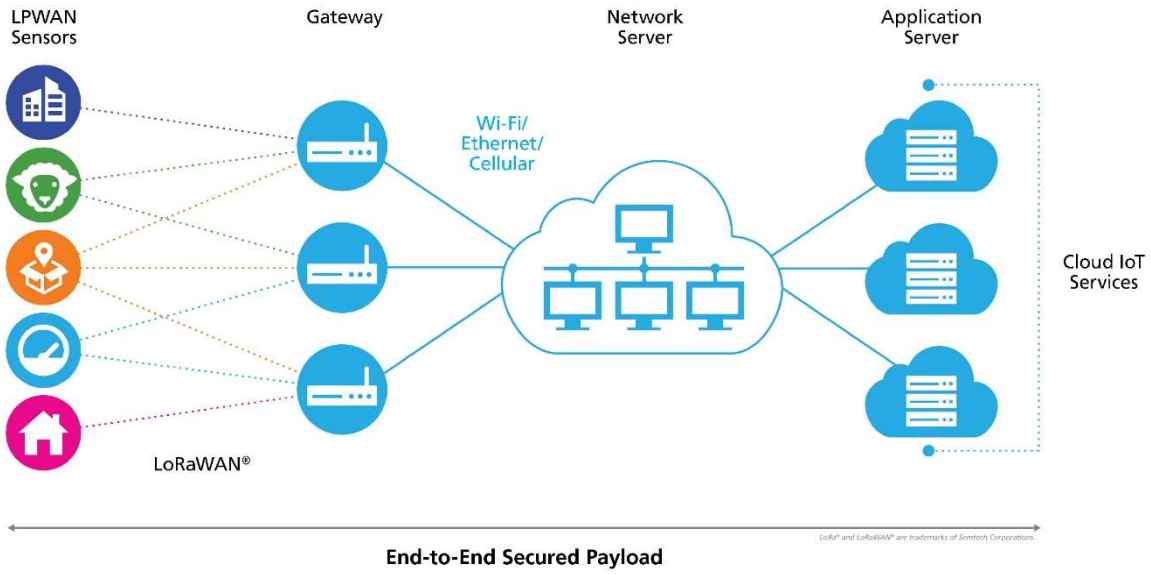

Figure 2: LoRa Technology enables the connection between remote point-of-use devices and LPWANs for delivery to analytics applications.
© Semtech Corporation. Reprinted with permission

Solving IoT Challenges

The open IoT LoRaWAN standard enables smart sensors, devices, and other objects to connect wirelessly via gateways that relay messages to a central network server. It is known for its bi-directional communication with end-to-end security, mobility, and localization services.

Additional advantages offered by the technology's wireless communication are significant compared to alternative options:

- Its robust long-range penetration reaches up to two miles in dense urban areas and offers an outdoor line of sight (LoS) up to 30 miles in remote rural landscapes.

- Semtech's LoRa Technology offers a 20dB link budget advantage, which significantly extends the range of any application while delivering the lowest current consumption to maximize battery life. In a LoRaWAN-based application battery-operated sensors can run up to 20 years, eliminating the need for the sensor power-source wiring that GSM, LTE, or Wi-Fi networks require.

- All LoRa-enabled devices are designed to comply with world-wide regulatory specifications, such as those of the FCC, ETSI, and ARIB, and with widely accepted standards including IEEE 802.15.4g and WMBus.

- LoRaWAN-based sensors and gateways (base stations) typically cost less than competing LPWAN technologies. There are no wireless spectrum license fees because

LoRaWAN-based technologies operate in the unlicensed spectrum. Cellular-based technologies operate in licensed spectrum and incur intellectual property (IP) royalties due to 3GPP heritage.

- The long range and energy efficiency of LoRa have enabled designers to eliminate the need for repeaters, thus reducing infrastructure cost, enhancing capacity, and scaling to support tens of thousands of nodes per gateway. Each gateway transmits the data securely to a cloud- or corporate-based server for analytics, empowering end users to make decisions based on hard real-time metrics rather than guesswork.

These performance characteristics have allowed LoRa to emerge as the go-to solution for networks demanding easy installation, low cost, and robust, secure communication in harsh environments.

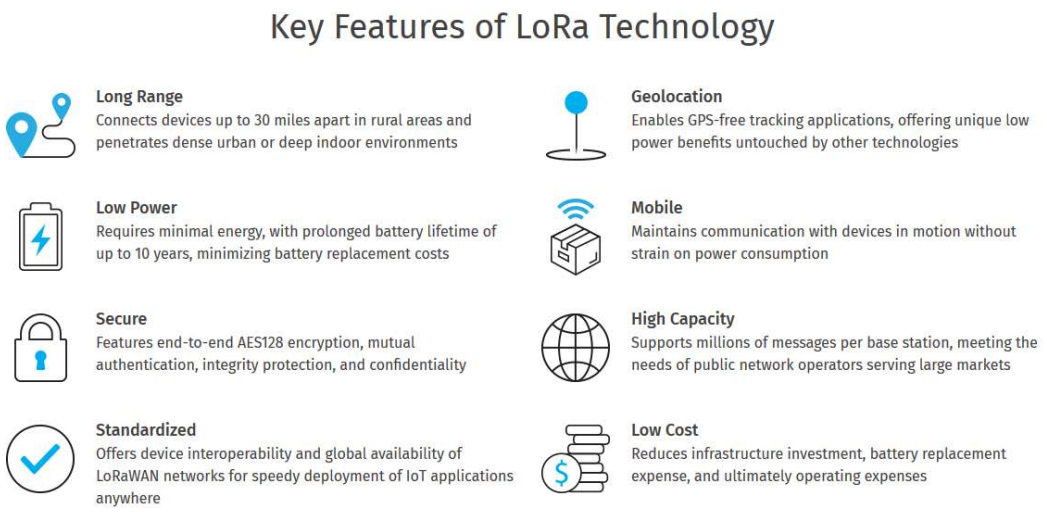

Figure 3: LoRa Technology's unique market advantages.

Simplifying the Future of Development and Accelerating Time to Market

Semtech's LoRa Technology enables innovators to overcome the age-old challenges that come with launching an IoT solution. The ecosystem supporting LoRa Technology is all-inclusive, providing end-to-end integration from silicon to services. As an open platform that operates on the unlicensed band, LoRa Technology is flexible for various business models to create profitability.

The standardized LoRaWAN protocol is interoperable, enabling solutions to scale, bundle, and evolve. LoRa Technology connects tomorrow's IoT solution today by simplifying the path for IoT innovators to bring products to market and demonstrate real-world use cases.

Semtech is rolling out a bundle of accelerators to streamline the process of developing, deploying, and managing LoRa-based IoT applications. The company is focused on understanding the needs of its customers – the solution providers and systems integrators who work every day to make the promise of the IoT a reality for their many customers across consumer, enterprise, and industrial segments – by providing easy to use, accessible tools, which they can use to more rapidly innovate.

LoRa Basics™: Basic code building blocks to assist solution developers in quickly realizing the ROI their customers want. The first of these building blocks, LoRa Basics™ Station (LoRaWAN gateway firmware), was announced and simultaneously released on GitHub in January 2019. Future areas of focus include device firmware, firmware updates over the air (FUOTA), and network performance analytics. To make accessing these tools easy, Semtech is revamping and evolving its LoRa Community™ into an open developer ecosystem portal to provide access to a comprehensive suite of developer training, technical resources, and community tools.

LoRa Cloud™ Services: Cloud services which deliver easy to use “ingredients” that solution providers can leverage to deliver value more rapidly with less development overhead. This service enables developers to quickly build IoT solutions which leverage multi-modal location capabilities (including Wi-Fi, GNSS, and LoRaWAN-based geolocation) without taking on the development complexity (and repetition) involved in building a solution from scratch. The service is designed from the outset to support flexibility in deployment options, cost effectiveness, and ease of use.

Modem-based, LoRa Hardware: Hardware platforms which simplify the deployment and management of IoT solutions. Semtech will be launching a new flexible modem-based hardware platform along with a cloud-based device provisioning and management service, which will dramatically simplify the full life cycle management of LoRa-based devices and accelerate both the development and deployment of secure fully-managed IoT solutions.

A Summary of LoRa Technology

LoRa Modulation: Semtech’s LoRa Technology is the physical (PHY) silicon layer, or wireless modulation, used to create the long range communication link.

Transceivers and End-Nodes: Transceivers configured with LoRa Technology are embedded into end-nodes, or sensor devices, designed for a multitude of industry applications.

Picocells and Gateways: Sensors capture and transmit data via LoRaWAN to gateways over distances near and far, indoor and outdoor, with minimal power requirement.

Network Server: Gateways send information via Wi-Fi, Ethernet, or cellular to the network server, which is responsible for network management functions like over-the-air activation, data

de-duplication, dynamic frame routing, adaptive rate control, traffic management, and administration.

Application Servers and Cloud IoT Services: Applications interpret the data collected by LoRa-enabled devices, applying techniques like machine learning and artificial intelligence to solve business problems for a Smarter Planet.

About Semtech's LoRa Technology

Semtech's LoRa devices and wireless radio frequency technology is a widely adopted long-range, low-power solution for IoT that gives telecom companies, IoT application makers and system integrators the feature set necessary to deploy low-cost, interoperable IoT networks, gateways, sensors, module products, and IoT services worldwide. IoT networks based on the LoRaWAN® specification have been deployed in 100 countries, and Semtech is a founding member of the LoRa Alliance™, the fastest growing IoT Alliance for Low Power Wide Area Network applications.

https://www.semtech.com/lora

CASE STUDIES

7.1 Real-Time Flood Monitoring with LoRaWAN

By Semtech

The Challenge

Noise pollution is a persistent problem for residents in urban environments. Recent studies have estimated that 9 in 10 people in major cities are exposed to noise levels exceeding international guidelines daily. The negative health effects of excess noise include disturbed sleep, hearing loss, cognitive disorders, and high blood pressure.

Municipal noise ordinances aim to reduce noise pollution, but assessments of noise and monitoring are performed infrequently and are primarily complaint-driven. In addition, noise monitoring is difficult to characterize and has traditionally been expensive. The City of Calgary set out to build a network of low-cost acoustic sensors to enable continuous monitoring in its urban environment.

The Solution

The Urban Alliance, a research partnership between the City of Calgary and the University of Calgary, was created to eliminate legal and financial red tape and coordinate the transfer of technology and research for the community's benefit. Dr. Henry Leung heads the Robotics and Sensor Networks Group in the Department of Electrical Computer Engineering at the University of Calgary. He collaborated with the Urban Alliance to build a LoRa-based sensor using Edge analytics to characterize noise and initiate a pilot use case on Calgary's LoRaWAN-based network.

"Our inspiration was the result of analyzing existing smart city noise monitoring applications. In New York City, they approached it using Wi-Fi and live streaming. This was rather expensive and took significant resources to reliably operate and sustain," said Leung. "We proposed developing our own solution without an electrical power supply to the device – just a battery."

Leung's team's first contribution to the project was the hardware design of the sensor. Their fabrication used low-power wide area radio transceivers to enable data transmission between the nodes and network server. The LoRa-based sensors are battery operated for ease of deployment, low power to limit network maintenance, and robust for continuous operation in extreme weather conditions. In addition they possess a limited amount of in-situ data processing.

The second contribution by the team was the development and testing of analytic algorithms allowing sensors to autonomously detect and classify acoustic events. The researchers will use

machine learning to distinguish among noise sources such as construction, traffic, gunshots, and music.

Improving Noise Monitoring in Urban Environments

To put these innovations to the test, several of the new LoRa-based sensors were placed at the Circle Carnival event at Shaw Millennium Park in September 2018. Installed at different locations around the park, the sensors were programmed to compute the average noise level every three minutes. When the noise level rose above 85dB, the sensors sent a warning packet through the LoRaWAN-based network. In the future, this feedback can be proactively provided to concert promoters to ensure noise restriction compliance.

From this testing it was observed that LoRa-based sensors were very easy to deploy with minimum infrastructure requirements. Testing also revealed the sensors were able to accurately detect noise thresholds without false alarm triggers.

The next evolution will be categorizing sounds such as trains, road noise, drag racing, gun shots, and construction and spatially correlating sounds over time and location. This data will help improve noise management and enforcement during public events by automatically alerting law enforcement when noise thresholds are exceeded, saving The City time and money.

About Semtech's LoRa Technology

Semtech's LoRa devices and wireless radio frequency technology is a widely adopted long-range, low-power solution for IoT that gives telecom companies, IoT application makers and system integrators the feature set necessary to deploy low-cost, interoperable IoT networks, gateways, sensors, module products, and IoT services worldwide. IoT networks based on the LoRaWAN® specification have been deployed in 100 countries, and Semtech is a founding member of the LoRa Alliance™, the fastest growing IoT Alliance for Low Power Wide Area Network applications.

https://www.semtech.com/lora

7.2 A story of the future: self-driving cars love parking sensors

While self-driving vehicles generate plenty of articles and discussions, one question deserves much greater consideration: Are those hundreds of autonomous vehicles expected to roam around the city center searching for a parking spot like humans do?

By Carlos Hernandez-Vaquero and Robert Bosch

Driverless Cars are just Around the Corner

By this time, probably most of us have briefly imagined a far-away future where most of the vehicles in the streets do not include a human as an active driver. It looks very much like science fiction, unreal or from a movie. It may also look scary and make us think twice about if this is ever going to work safely enough in real life conditions.

There have been different predictions about the speed at which driverless cars will become mainstream. The Society of Automotive Engineers (SAE) has defined several levels of driving automation from SAE level 0 (not automated) to SAE level 5 (fully automated). Figure 1.

SAE level	Name	Narrative Definition	Execution of Steering and Acceleration/ Deceleration	*Monitoring* of Driving Environment	Fallback Performance of *Dynamic Driving Task*	System Capability *(Driving Modes)*
***Human driver* monitors the driving environment**						
0	**No Automation**	the full-time performance by the *human driver* of all aspects of the *dynamic driving task*, even when enhanced by warning or intervention systems	Human driver	Human driver	Human driver	n/a
1	**Driver Assistance**	the *driving mode*-specific execution by a driver assistance system of either steering or acceleration/deceleration using information about the driving environment and with the expectation that the *human driver* perform all remaining aspects of the *dynamic driving task*	Human driver and system	Human driver	Human driver	Some driving modes
2	**Partial Automation**	the *driving mode*-specific execution by one or more driver assistance systems of both steering and acceleration/ deceleration using information about the driving environment and with the expectation that the *human driver* perform all remaining aspects of the *dynamic driving task*	**System**	Human driver	Human driver	Some driving modes
***Automated driving system* ("system") monitors the driving environment**						
3	**Conditional Automation**	the *driving mode*-specific performance by an *automated driving system* of all aspects of the dynamic driving task with the expectation that the *human driver* will respond appropriately to a *request to intervene*	System	**System**	Human driver	Some driving modes
4	**High Automation**	the *driving mode*-specific performance by an automated driving system of all aspects of the *dynamic driving task*, even if a *human driver* does not respond appropriately to a *request to intervene*	System	System	**System**	Some driving modes
5	**Full Automation**	the full-time performance by an *automated driving system* of all aspects of the *dynamic driving task* under all roadway and environmental conditions that can be managed by a *human driver*	System	System	System	**All driving modes**

Figure 1: The five levels of autonomous driving as defined by SAE International.

While some levels of automation are already present in the cars we can buy today, this is currently only SAE level 2, including the famous Autopilot in the Tesla cars. This automation level might temporarily allow us to take our hands off the steering wheel, but we still need to be looking at the road and ready to take control when required.

Some car manufacturers, the first being Audi with the latest A8, are now announcing the availability of SAE level 3 features in their cars, which allow the driver to do something else while driving, like reading the news, but require the driver be ready to take control when necessary.

The magic however starts at SAE levels 4 and 5, where the driver is not required partially or completely anymore. He or she may be sleeping in the car or simply not there. This opens a complete new world of dilemmas and opportunities.

Most of the forecasts place the marker from 2020 to 2025 for massive commercialization of some SAE level 4 features, while reaching complete automation around 2030.

Let’s go to the facts.

Waymo, previously called the Google Self-Driving Car Project, recently reported driving over 10 million miles with its self-driving cars. They also released a product called Waymo One, which although still only available for selected customers, is being charged for in manner similar to a taxi service. A human is in the driver’s seat for security reasons, but he is generally not doing anything, just looking around.
As surprising as it sounds, this is already happening: autonomous cars transporting customers to their destinations for a price, slowly but safely.

On the European side, Daimler and Bosch have been cautiously working together to deliver a highly automated (SAE level 4) and driverless (SAE level 5) robo-taxi experience from the beginning of the 2020’s. That sounds like a lot of time, but in reality, this may be no more than one or two years from now for the pilot cities.

Expectations in the Daimler and Bosch case go beyond a service which is restricted to a handful of customers, features slow moving cars, and requires a human driver’s presence. Being able to order a taxi without a driver in every country with medium to high labor costs is not science fiction anymore.
We should still be realistic. The replacement of millions of non-autonomous cars by autonomous cars will take several decades. During the next 20 or 30 years there will be a mix of “old” and “new” vehicles. However, it is reasonable to think that inhabitants of several cities in developed countries will start living with some type of self-driving cars before their next local elections.

The Impact of Autonomous Cars on Daily Life

As soon as there is a relevant amount of highly automated (at least SAE level 4) vehicles in the same city, new and innovative services and jobs will begin to be in demand. The self-driving cars will generate unforeseen dilemmas, which in one way or another will transform our cities and society.

Unfortunately, the sudden extinction of thousands of jobs related to transportation is unavoidable, and we will see many frustrated citizens taking to the streets to press their governments to preserve the status quo. In addition, if our city administrators do not take the right actions, we may have the risk of flooding certain areas in big cities with a complete army of machines running on wheels. For example, after a football match, they may try to pick up thousands of spectators at the same time in the same place or congest the whole city center while searching for a place to park.

One of today's clearest advantages for self-driving car users is being dropped off at the exact place they want to go while the car takes care of finding a parking place by itself. Drivers who arrive in non-autonomous cars might spend would spend around 20 minutes searching for a parking space in a typical city center while contributing to increased traffic and therefore slowing other drivers' search for a parking space.
What's beneficial at the individual level—having the car find parking—could be a big new problem for cities as a whole and mean new headaches for city administrators who have been fighting for decades to reduce the traffic in our city centers.

A clear example is the Park and Ride (P+R) system, which has been developed and promoted in many cities as a solution for reducing city center congestion. Who will now park in the suburbs and get on public transport when our cars can drop us at our final destination and go by themselves to find a parking spot, even if that spot is back in the suburbs?
P+R has been relatively successful in big cities because we don't want to waste our time on top of having to pay for parking while we go to work, do shopping, meet a friend, or do anything else.

With self-driving cars, we no longer need to spend our valuable time searching for parking in addition to paying high parking fees. Our smart vehicle is now able to spend its non-valuable time searching for the most convenient parking space, with "convenient" most likely defined as a balance between the cheapest and the nearest. In some cases, for a car to simply continue driving around until its passengers are ready to return may be the best choice.

As if our cities would not have enough problems by now, we may see a city center flooded by autonomous cars driving in circles, marking time until summoned by their travelers. For them, joining a traffic jam might even be the best possible scenario because this could potentially consume less energy.

We can imagine a Hollywood movie whose plot includes the hero trying to get across a city fast only to be challenged by a massive artificial traffic jam of autonomous vehicles created by the desire to avoid both parking fees and high energy consumption while they wait for their travelers.

Looking at the future, anticipating the changes, and taking the right decisions at the right time are the tasks of our politicians and entrepreneurs. Those who are brave and technologically ready will embrace the changes and benefit from the opportunities created with these game-changing technologies. The magnitude of the impact that will be produced by autonomous cars is comparable to only a few other events, which seem to happen once or twice in a century.

In any case, if the city administrators are open to accept and take part in technology advancements instead of opposing them, some of the new dilemmas may have a retro-fit solution, which does not require excessive investments in infrastructure.

Making Opportunities a Reality

Many of the people who live in the city center use a long-term underground private parking solution, because on-street public parking is scarce and usually not free. Renting or owning an underground parking space is in the long run usually much more convenient and cheaper than “pay as you go” solutions like public parking areas.

However, the renter or owner of such private underground parking solutions must still invest a significant amount of money for that added convenience, a luxury in some cases. But they also travel with their car and therefore leave their parking space empty during long hours every week.

In the meantime, in the streets near these “invisible” empty parking spaces, dozens of cars roam around searching for the nearest available parking space or finally deciding to drive several hundred meters to an expensive public parking lot. In the worst case, during special events, there might be no available public parking lots at all in the nearest couple of kilometers.

How is it possible that these expensive resources are left underutilized while there is high demand for them? There are, in fact, several startups, which are innovating to create a collaborative economy for parking spaces, renting private parking spaces while they are not in use. This is currently happening with relative success, but the demand for such a solution may explode as autonomous car sales take off.
In 2019, these collaborative economy solutions force the parking space requester to manually search in a mobile app for available short-term parking spaces, which also requires a second passenger to support the driver, because he or she still needs to drive. In addition, the lack of standardized solutions puts up barriers to adoption by adding several layers of complexity to what should be straightforward actions. That would include such things as managing access to the private parking space and detecting free parking spaces.

Another possible solution to parking space scarcity at the time driverless cars go mainstream involves centralized smart traffic control. In this case, the city administrators may decide to allow autonomous cars to park in a second or third row in a street that does not require all its lanes due to temporary low traffic density.

If any blocked car in the first row is summoned, the blocking cars will simply make space, like in Tetris, for that car to leave. Next, the remaining cars will arrange themselves so as to reduce the number of movements required to allow the next departing car out. The mechanism by which they arrange themselves may depend on each car's expected utilization time.

Of course, the above concept will not work if any of those cars in the second or third row are not fully automated and require a human to act as the driver. In such a case, the city administration may need to summon an autonomous tow truck to remove the offending car.

A Parking Lot Sensor Solution for Making "Invisible" Parking Spaces Usable

While much of the media attention has been focused on the development of self-driving vehicles, some companies are working in the background to start covering those vehicles' future demands for specific products and services, which all of a sudden, may be urgently needed.

One example is the Bosch Parking Lot Sensor (Figure 2), which is glued on the ground of a parking space and reports the current status— either busy or free—with a smart self-learning algorithm. This sensor uses LoRaWAN and is an enabler to make usable all those "invisible" empty parking spaces in private areas.
Although traditional cellular networks are already publicly available in many cities, the conditions required for a parking space sensor dovetail perfectly with LoRaWAN strengths. Underground parking areas, or even the attenuation produced by the car parked on top of the sensor in the street, generally require private deployments where traditional networks are too weak or too power hungry.

Parking sensors speak machine language, allowing the automation of future solutions for increased parking utilization. The access management system to the private parking areas could be simply solved by applying a concept similar to Bosch Perfectly Keyless to the parking physical barrier. Bosch Perfectly Keyless is a digital key on a mobile phone which allows secure access and start of a car to its owner through a mobile app. In a similar way, autonomous cars could be able to authenticate and interact with physical elements without the need of the current physical devices like remote controllers or keys.

Only the massive deployment of parking lot sensors and a digital platform to connect autonomous cars with free parking spaces can remove the barriers, allowing the success of a collaborative economy for the parking spaces.

The collaborative economy concept has already demonstrated its value, as with Airbnb, Couchsurfing, Uber, Car2Go, Kickstarter, Amazon Mechanical Turk (MTurk) or eBay, to cite a few examples. Sooner or later obtaining a parking space will also be part of the collaborative economy.

Today, fundamental reasons for the deployment of parking lot sensors, aside from any autonomous vehicle needs, already exist: e-vehicle charging stations, monitoring of critical emergency vehicle exits, tracking utilization per parking space, etc.

Nevertheless, once the self-driving cars start conquering the streets, the parking sensors will be simply considered a basic ingredient for the smart city, in combination with smart traffic control, waste management, automated planned emergency response to catastrophic events, and many others to come.

Figure 2: Bosch Parking Lot Sensor

About the Author

Carlos Hernandez-Vaquero is a Software Project Manager and Product Owner of several Bosch IoT products in the area of Stuttgart, Germany. He received his Telecommunications Engineering degree in the University of Oviedo (Spain) and his M.Sc. degree in Wireless Communication Systems in the University of Aalborg (Denmark).

He is especially proud of having published more than 20 mobile apps and the release to the market of the Bosch Transport Data Logger. He has co-authored several patent applications and contributed to the standardization of the IO-Link technology. He is a certified Scrum Product Owner and CompTIA Security+ expert.

Carlos has worked in different projects at Intel and Omron Corporation, but he is now obsessed with the Bosch Parking Lot Sensor, where he believes a revolution is coming hand in hand with the Smart Cities.

With his origins in a beautiful town by the beach in Spain called Salinas, he moved in 2010 to Denmark and later to Germany to try to learn from the best engineers he could ever imagine.
Carlos is continuously expanding his horizons to other domains and he is now preparing in his free time the AWS Solutions Architect and PMP certifications.

https://www.bosch.com/

7.3 LoRaWAN Technology: Transforming Golf Courses by Monitoring Pace of Play

By Semtech

The Challenge

In 2015, the City of Calgary celebrated 100 years of municipal golf at its first public course. Shaganappi Point began as a bare-bones 18-hole golf course serving a community of 80,000 Calgarians. Calgary's population has grown to more than one million, and Shaganappi Point has expanded to 27 holes, a 44-stall driving range, and a club house. During its annual season, from April to November, golfers play an average of 90,000 to 100,000 rounds of golf.

The staff at Shaganappi Point wanted to measure the amount of time it takes each guest to complete a round of golf on its course. They wanted to know if this actionable information could elevate customer experience and maximize revenue. Previous issues with "slow play" during peak periods had reduced customer satisfaction and may have hurt customer retention.

The Solution

Calgary Recreation worked with Information Technology at The City of Calgary to embed golf carts with tracking sensors to provide insight about factors affecting the pace of play. For an initial proof of concept, small units of LoRa-based-sensors were placed underneath the seats of [how many?] golf carts. A LoRaWAN-based network transmitted the data collected by those sensors to a TEKTELIC Communications KONA Mega IoT Gateway placed on a 100-meter City-owned radio tower seven kilometers from the area.

The gateway was connected to a TEKTELIC network server, with GlobalSat supplying the GPS sensors. Custom software, developed by SensorUp, provided the necessary real-time and historical course usage data on a web-based dashboard that could be viewed on a computer or mobile device.

Enabling New Efficiencies and Cost Savings

Being able to obtain real-time golf cart location information enhanced facility operations and improved guests' experiences. As pace of play anomalies were detected, course marshals were dispatched to support golfers needing assistance. Shaganappi Point is also using sensor data to determine accurate tee time intervals, make course adjustments, and better predict revenue impacts.

The City of Calgary is one of the first cities in North America to build a municipality-owned, carrier-grade LoRaWAN-based network. LoRaWAN technology has been in place for over one year and is a part of a City Network of Things (CNoT) platform created by the City of Calgary's

Innovation & Collaboration team. Calgary envisions its CNoT will be used by many of the City's 32 business units to eventually connect tens of thousands of sensors.

https://www.semtech.com/lora