



THE DEFINITIVE GUIDE

THE INTERNET OF THINGS FOR BUSINESS

3RD EDITION



FOREWORD

The Internet of Things is hitting its stride. It is a critical enabler of digital transformation efforts happening around the globe. For consumers, businesses, and governments, IoT is playing an increasing role in how we consume information to make decisions and how we interact with the world around us. More than that, IoT also is changing how we personally experience our world—our daily interactions with our connected car, smart homes, and connected wearables make us digitally linked to the physical world in a way that would seem like magic to people living just 100 years ago.

We've journeyed from a world of M2M (Machine-to-Machine), which was little more than connected endpoints providing descriptive data about a "thing's" behavior, to an IoT world of intricate and complex endpoints capable of capturing and transmitting predictive information.

Because of this, Cloud has become a critical enabler of many successful IoT deployments, yet the need to enable processing and computational capabilities at the edge of the network is becoming increasingly important. IDC expects that in the next phase of the IoT, more than 50% of data will be processed at

the edge. Data captured will become more prescriptive and intelligent—allowing proactive actions to be taken based on these insights. This guide examines the increasing use of data as a driver of business, as well as other critical aspects of IoT.

Successfully planning, deploying, and scaling IoT solutions is a challenge for businesses in every industry, and the goal of this book is to provide entry and insight into the IoT world. It also serves as a valuable guide for businesses who want to leverage IoT to expand revenue, stay relevant, and advance business objectives in the connected economy.

Whether you want to attempt initial entry into the IoT-sphere, or expand your existing deployments, this book can help you meet your goals, providing deep understanding into all aspects of IoT. From concept to lifecycle management, or platform selection to data analytics, *The Internet of Things for Business* (3rd ed.) guides each step of the way.

Carrie MacGillivray

Group Vice President & Global IoT Lead
IDC



The Definitive Guide | The Internet of Things for Business, 3rd Edition

By Syed Zaeem Hosain, CTO, Aeris

Aeris® and AerPort™ are the trademarks and / or registered trademarks of Aeris Communications, Inc. All third-party trademarks are the property of their respective owners.

Copyright © 2018 Aeris Communications, Inc. All rights reserved. No part of this book may be used or reproduced in any manner whatsoever without the explicit permission of the publisher.

Third Edition: August 2018

Editor: Carmi Brandis

Creative Director: Nick Waschezyn

Book Design: Delin Design

Subject Matter Experts:

Michelle Avary
Stephen Blackburn
Racquel Brown
Yixiang Chen
Christina Richards
Narendra Sharma
Evan Whitelock

Project Management:

Bob Heckmann
Christina Richards

For additional information about this book, contact Aeris Communications, Inc. or visit www.aeris.com.

Also, follow us on Twitter [@AerisM2M](https://twitter.com/AerisM2M) to learn how we can inspire you to create new business models and to participate in the revolution of the Internet of Things.

U.S. Global Headquarters

Aeris Communications, Inc.
1745 Technology Drive
Suite 700
San Jose, CA 95110

Chicago

435 N. LaSalle Drive
Chicago, IL 60654

Europe

Unit 1, Blake House
Manor Park, Basingstoke Road
Reading RG2 0NA UK

India

6th Floor, Tower One
Okaya Business Center
Yoganand Marg, Sec-62
Noida 201 309
Uttar Pradesh



CONTENTS

1

1 WHAT IS THE INTERNET OF THINGS?

- 3 Machine-To-Machine (M2M)
- 4 Internet of Things (IoT)
- 6 Impact for Businesses and Consumers
- 6 Consumer IoT Applications
- 8 Enterprise IoT Applications
- 10 Using IoT for a Better World
- 12 An Intro to the Guide to IoT For Business

2

14 THE FUTURE OF PLATFORMS

- 15 IoT Platform Layers and Services
- 19 From Many to One
- 21 Advantages of a Single vs Multi-Platform Deployment
- 22 Platform Requirements

3

25 IoT CONNECTIVITY: TYPES AND CHOICES

- 27 Basic Internet Concepts
- 28 Choice of Connectivity
- 32 Proprietary Protocols
- 34 Standardized Protocols
- 37 Types of Cellular Technologies
- 44 Cellular Fallback
- 45 How To Determine Location
- 47 Global Positioning System

4

52 CONNECTIVITY MANAGEMENT PLATFORMS

- 54 What is a Connectivity Management Platform?
- 55 The Difficulties of Managing IoT Connectivity
- 56 Why Businesses Need Connectivity Management Platforms
- 59 Essential Connectivity Management Platform Features
- 61 CMPs are Integral to the IoT Environment

CONTENTS

5

63 IoT SENSORS AND DATA COLLECTION

- 65 What is a Sensor?
- 66 Sensor Types
- 72 Conversion to Digital Data
- 76 Calibration and Linearization

6

78 IoT ANALYTICS

- 80 IoT Data and Analytics
- 81 Types of Analytics
- 85 Future of Analytics

7

87 SCHEDULING, ENCODING, AND PROCESSING

- 89 Data Transmission Schedules
- 91 UDP or TCP
- 93 Content Encoding / Transport Protocols
- 98 Gateways
- 98 Application Servers
- 99 Cloud Computing
- 100 Fog Computing

8

102 IMPLEMENTING AN IoT SOLUTION

- 104 Supply Chain Management
- 104 Cellular Operator Selection
- 106 Cloud System Selection
- 107 Platform Selection
- 107 Network Operator Service Level Agreement
- 108 Device Certification
- 109 Considerations
- 111 Application Communications Call Flow
- 112 Customer Support Process

9

113 IoT SCALABILITY AND ALTERNATIVE TECHNOLOGIES

- 117 What is Scalability
- 119 End-of-Life Management
- 120 Scalability and Connectivity

CONTENTS

10

128 SECURITY, PRIVACY, AND THE INTERNET OF THINGS

- 130 Privacy and Security
- 133 International Data Transport
- 133 Security Objectives
- 135 Security Issues for IoT
- 137 Risk Management and Assessing Impact of Breaches
- 141 Encryption as an IoT Tool
- 142 Choice of Encryption Algorithm

11

143 IoT USE CASES

- 145 Renewable Solar Energy
- 146 Automotive
- 149 Healthcare
- 152 Smart Cities
- 154 Financial / Insurance

12

156 THE FUTURE OF THE INTERNET OF THINGS

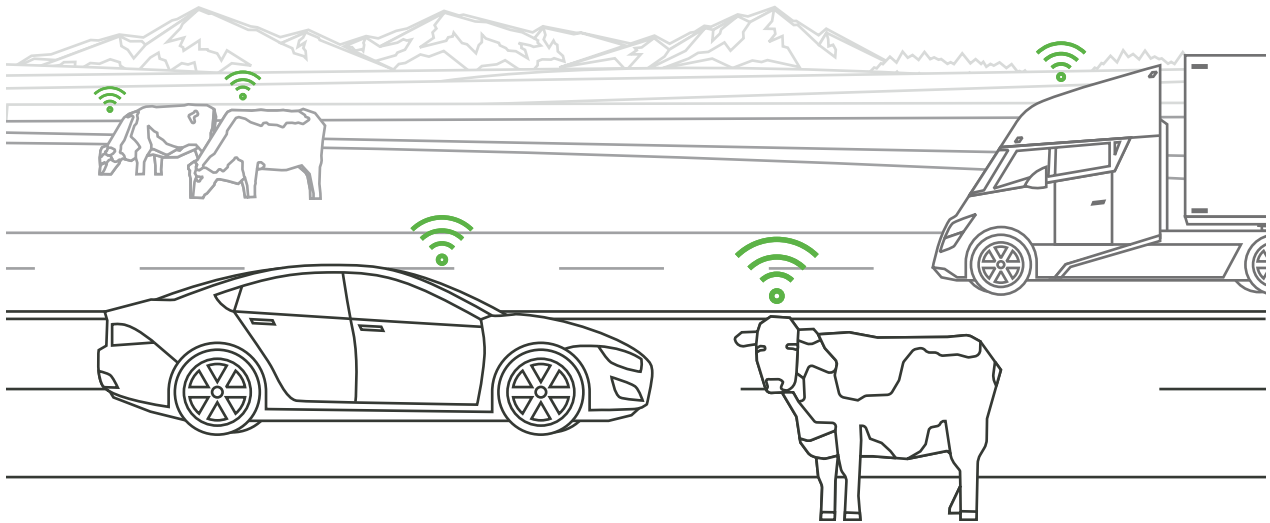
- 158 IoT Will Positively Affect All Markets
- 158 IoT Will Come First
- 159 Homes Will Get Smarter, And More Connected
- 160 Enterprises Will Spend More on IoT
- 160 IoT Standards Will Improve
- 163 Security Concerns Will Continue
- 164 Over-The-Air Updates Will Become the Norm
- 165 Privacy Concerns and Government Regulations
- 166 IoT Value Realized Through Data Analytics
- 166 The Future Is Now

167 DIRECTORY OF IoT TERMS

- 168 Acronyms
- 178 Glossary

WHAT IS THE INTERNET OF THINGS?

3	MACHINE-TO-MACHINE (M2M)
4	INTERNET OF THINGS (IoT)
6	IMPACT FOR BUSINESSES AND CONSUMERS
6	CONSUMER IoT APPLICATIONS
8	ENTERPRISE IoT APPLICATIONS
10	USING IoT FOR A BETTER WORLD
12	AN INTRO TO THE GUIDE TO IoT FOR BUSINESS



WHAT IS THE INTERNET OF THINGS?

The Internet of Things envisions a world where both ordinary and exotic devices are connected wirelessly to the internet and to each other. This means devices that do not already have a network connection may have one added in the future, when it is logical and appropriate to do so.

For example, an IoT device could be a temperature gauge, a location sensor, a device measuring humidity, or a vibration detector. One or all of these sensors then could be attached to manufacturing machinery, and the data transmitted would help a business track the machine's operations. This data could track required maintenance, improve production efficiencies, reduce downtime, increase safety, and more. Plus, IoT devices can provide information on the ambient environment of the manufacturing space, such as the temperature, pollution, and other conditions near the machinery, which can be particularly relevant for remote installations.

Most IoT projects are motivated by a need to reduce operating costs or increase revenue. Occasionally, legislation compels companies to deploy IoT applications that support a new law's data needs. Mobility is an obvious factor driving cellular adoption in markets like transportation. Desire for competitive features will inspire IoT applications in consumer high-tech. But whatever the specific purpose, connected IoT devices can give your business the data and information needed to streamline workflows, predict necessary maintenance, analyze usage patterns, automate manufacturing, and more.

However, the very term “Internet of Things,” coined by British entrepreneur Kevin Ashton in 1999, may no longer apply in its original form.

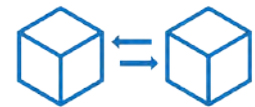
Some clarity is needed to differentiate Internet of Things from Machine-to-Machine technologies. Following are definitions and examples of these overlapping technologies, and how they have evolved over time to improve business efficiencies, produce multiple new revenue streams, or simply enrich the quality of life.

MACHINE-TO-MACHINE (M2M)

By many indicators, M2M seems a lot like IoT. The difference, however, is that M2M is a solution that optimizes existing operations functionality through automation, while IoT transforms the functionality into new business capabilities via analytics. Hence, M2M can be thought of as a sub-set of IoT.

Though initially not built as a sub-set of anything, M2M technologies represent closed, point-to-point communications between machines or between machine and management systems, without the need of human intervention. M2M devices, enabling bidirectional remote monitoring and transfer of data, consist of a sensor or an RFID tag and a communication module. Machine-to-machine devices, as the industrial precursors to the IoT, can include items ranging from in-house / in-office machinery, such as printers or scanners to manufacturing equipment, including heavy machinery. But don't assume that the IoT will replace M2M. Predictions show that cumulative M2M connections will grow from 995 million in 2014 to a projected 2.7 billion connections by 2018.

M2M use cases include telemetry; traffic control; security; tracking and tracing; machinery maintenance and control; metering; manufacturing and facility management; as well as a multitude of additional applications.



INTERNET OF THINGS (IoT)

The Internet of Things goes beyond the scope of M2M, encompassing and surpassing it in functionality by adding devices and electronic equipment with embedded sensors, control systems, and processors that enable communication across a multi-node, open network of objects.

An Internet of Things ‘thing’ can refer to a connected medical device, a biochip transponder for livestock, a solar panel, a connected automobile with sensors that alert the driver to a myriad of possible issues (fuel, tire pressure, need for maintenance, and more), or any object, outfitted with sensors, that has the ability to gather and transfer data over a network.

The meaning and application of the term IoT will continue to evolve as new connected technologies emerge. For many, IoT means connecting parts of the supply chain, increasing proficiency and outcomes, and providing indicators about product environments. For others, IoT is about life-changing insights via wearables, medical adherence, or household security. The possibilities just keep growing.

IoT use cases are widespread, limited only by our ability to connect certain devices. But that, too, is changing rapidly. IoT applications already in production include connected cars, smart cities (water / gas meters, lighting, traffic / parking, waste management, and more), patient monitoring / medical adherence, wearables, agriculture, and energy, with more uses seeing daylight with each passing week. IoT use cases in the future will be limited only by our own imagination as, eventually, all other barriers will fall.

Connected for Data

The Internet of Things is the next logical step in the story of a connected world. At the heart of the IoT is data—the ability to collect it, analyze it, and react to it, so as to create new revenue streams, new value. The IoT combines the technologies found in M2M and earlier data telemetry terms and expands them with an even greater accumulation of data and inferences.



The IoT is comprised of four key elements:



People—Using end-nodes connected to the internet to share information and activities. Examples include social networks, health, and fitness sensors, among others.



Things—Physical sensors, devices, actuators, and other items generating data or receiving information from other sources. Examples include smart thermostats and home automation gadgets.

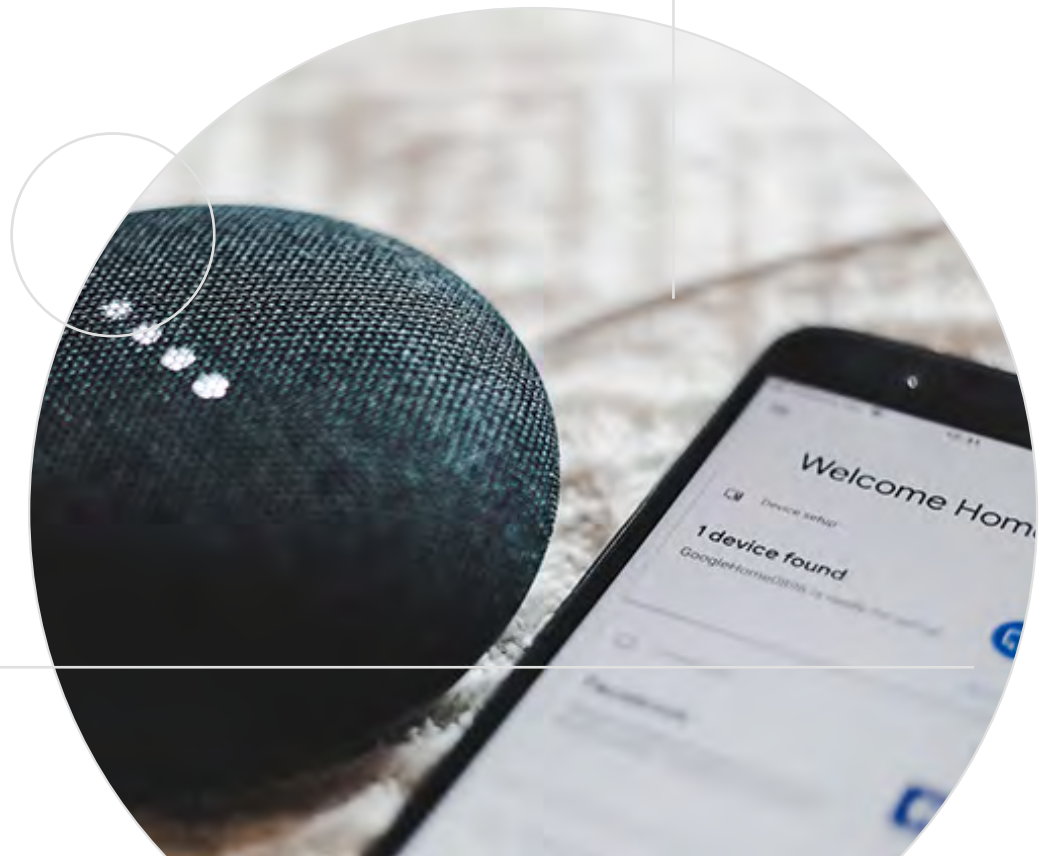


Data—Raw data analyzed and processed into useful information to enable intelligent decisions and control mechanisms. An example includes temperature logs converted into an average number of high-temperature hours per day to evaluate energy requirements.



Processes—Leveraging connectivity among data, things, and people to add value. An example of this includes the use of smart fitness devices and social networks to advertise relevant healthcare offerings to prospective customers.

The IoT establishes an end-to-end ecosystem, including technologies, processes, and concepts employed across all connectivity use cases.



IMPACT FOR BUSINESSES AND CONSUMERS

The concepts of the IoT and M2M are inherently subject to the confusion surrounding limitations associated with meanings, use cases, and adoption.

While there are not yet comprehensive standards and regulations for IoT from appropriate authorities (such as 3GPP¹), these concepts will continue to evolve in response to technology innovation, changing consumer trends, and varied marketing tactics. Businesses evaluating the promise and potential of connectivity offerings will, therefore, have to dig into the specifics of each situation instead of establishing conclusions based solely on the proposed labels of IoT or M2M.

As new as the Internet of Things may seem, many network-connected devices already are in use all around us. You probably have heard of connected homes or the smart grid—these are just a few of the IoT systems aimed at both everyday consumers and large-scale enterprises.

IoT innovation is taking place in a wide range of industries, locations, and types of business. IoT creativity will be unlimited, as the technology largely exists—although it may not be readily available everywhere as of yet. Or as William Gibson so famously stated, “The future is here. It’s just not widely distributed yet.”



CONSUMER IoT APPLICATIONS

While the focus of this book is on business uses for IoT technology, seeing how it applies to consumer devices is relevant for a sense of scale and direct application in everyday lives.

IoT devices let individuals control their own network-connected devices from their smartphones or wearables or get information about their status from a webpage.



¹ The 3rd Generation Partnership Project (3GPP) is an association of seven telecommunications standard development organizations (ARIB, ATIS, CCSA, ETSI, TSDSI, TTA, TTC), providing standardization oversight for cellular telecommunications network technologies.

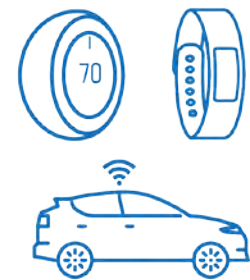
The most popular consumer IoT devices typically are found in three major categories. The first is the connected home, which includes the smart thermostat, intelligent lights, connected appliances, and smart door locks. Next, wearables dominate the consumer market with the smartwatch, activity / fitness tracker, and smart glasses. Finally, the connected car rounds out the consumer categories with remote car controls, trip navigation, and vehicle diagnostics.

Here are a couple of examples of popular consumer IoT applications:

The Nest thermostat arguably is one of the most well-known of the products in this category. Nest, currently owned by Google, provides a Wi-Fi-connected thermostat that is capable of learning a person's activities and setting room temperature based on those preferences. The idea behind the product is to always keep a home comfortable while boosting energy efficiency. The Nest thermostat can be integrated with automated IoT lighting, security systems, and other tools, thereby making the long-imagined connected home more of a reality.

Internet-connected fitness trackers, such as Fitbit and smartwatches like the Apple Watch, do everything from act as pedometers to sleep alarms to personal coaches. These devices are part of a "quantified self" movement that started in the mid-2000s to gain greater personal understanding through data and technology. Devotees feel that these wearables help to achieve health goals, and they even are used by businesses as part of employee wellness programs to incentivize fitness and, potentially, reduce health insurance premiums.

The connected car is one IoT application that has witnessed a large increase in features. Many cars come equipped with systems that gather data within the vehicle and can transmit them for a variety of applications. After-market devices capture sensor data using the vehicle's on-board diagnostic port (OBD-II) for cars built since 1996. Examples include automatic notification of crashes, notification of speeding, and safety alerts. Additionally, concierge features provided by automakers or apps alert the driver of the best time to leave for a prompt arrival to an appointment or sending text message alerts to friends or business associates to alert them of arrival times. Users also can unlock their cars, check the status of batteries on electric vehicles, find the location of the car in a parking lot, or remotely activate the climate control systems. As time passes, we expect an increasing number of applications, including a truly self-driving or autonomous car, to be made possible by IoT technology.



ENTERPRISE IoT APPLICATIONS

To date, most industrial uses of IoT have been for preventive maintenance. These applications detect when a machine has variations in vibration, temperature, speed, or other metrics so as to signal that they might require maintenance.

But using IoT for preventative maintenance was just a start. This didn't fully tap into the ability of network-connected devices to talk to each other, thus letting them work together. For example, a business could use a central monitoring hub, or even an engineer with a smartphone, to reach out to the machine and make changes on the device, or deliver new instructions. More and more enterprises are realizing that these communications can create greater efficiencies and reduce production costs far beyond IoT systems aimed at simple maintenance functions.

The fleet industry was one of the earliest to adopt IoT because of its many benefits. IoT-enabled trucks, ships, and vans can be tracked and managed in a more efficient manner, thereby allowing visibility across the transportation ecosystem. Fleet telematics allow the exchange of information between a commercial vehicle fleet and a central dispatching office. Now, the physical health of a vehicle can be checked at a fraction of the cost and in real time. Additionally, GNSS-based tracking can guide a vehicle to its destination in the most efficient manner and allow the central office to optimize the dispatch of its fleet more effectively. Some of the leaders in the fleet management space include PeopleNet and Omnitrac.

Here are examples of other industries with interesting IoT enterprise applications that currently are deployed:

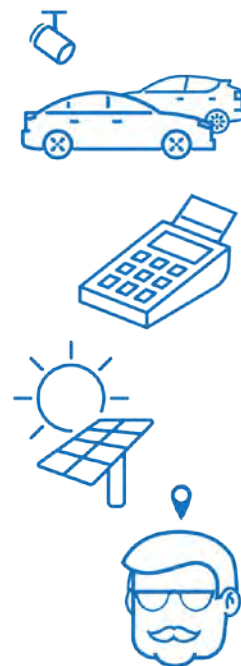
Point of Sale: Acceptacard is a provider of dedicated card-processing solutions for UK businesses. Its mobile point-of-sale (POS) terminal is a breakthrough from what typically is provided by the banking industry in that there are no multi-year contracts with expensive terminals. The company caters to a broad range of business types—from startups to established companies looking to reduce costs or for greater payment functionality, flexibility, and control. Its mobile payment solution is a terminal-independent solution with reliable connectivity service, regardless of the location, and is designed for businesses that want a payment solution on a self-service basis with online access.



Ride Sharing: Uber is fast becoming synonymous with IoT ride-sharing services, regardless of the location or the mode of on-road transportation. Combining cars, shared carpools, decentralized scooters and bikes, car rentals, and even public transportation, Uber tries to find the most efficient way to get you from here to there. Leveraging smartphones as a basis for its business, Uber connects drivers with those needing conveyance. By relating passenger locations with its drivers, Uber can route services optimally to maximize results. And it doesn't end with just driver services. Uber continues to invest in self-driving technologies, acknowledging the near-future direction of vehicles. Additionally, Uber uses other modes of transportation for that 'last mile' of transit. Bikes, scooters, and Uber Pools all can alleviate traffic congestion while providing mobility services to its clientele. In fact, Uber is continuing to invest in new people-moving technology, including dockless scooters and bikes (covering 70+ markets with more than 35,000 scooters in service nationally). And within some cities, Uber has even partnered with public transit, selling tickets via its own app. For inventive companies such as Uber, the IoT is their highway to success.

Solar Energy: BBOXX designs, manufactures, distributes, and finances innovative plug-and play, off-grid solar powered systems to improve access to energy across Africa and the developing world. Because of the importance of sustainable energy, BBOXX aims to provide 20 million people with electricity by 2020. Its core products include a range of solar powered battery boxes that sit in a home and allow users to power small appliances, such as lights, mobile phones, refrigerators, or computers. BBOXX has more than 80,000 systems deployed so far across China, UK, and East Africa.

Healthcare: SimplyHome designs and installs wireless technology products and related home care-focused services. The company is known for its highly customizable systems that are tailored to meet each customer's specific needs. Its systems proactively alert patients and caregivers to changes in behavioral patterns by communicating with multiple sensors to observe activities of daily living. Its products and services range from voice-activated environmental controls, personal emergency response systems, GPS watches, motion sensors, stove monitors, and virtual care management. Text, email, or phone alerts can be generated by a single event, an intersection of multiple events, or by inactivity. The SimplyHome system helps residents remain independent with environmental controls that operate beds, lights, TVs, doors, and more via tablet or voice activation.



USING IoT FOR A BETTER WORLD

While many enterprises are using IoT technology to make money, nonprofit organizations and non-governmental organizations (NGOs) are showing how IoT can be used to make the planet a more habitable place and improve people's quality of life.

Aeris provides global connectivity services and solutions across multiple carriers and multiple technologies for social impact enterprises working in some of the world's most challenging environments.

Here are a couple of examples.


SweetSense is an organization that has teamed with governments and NGOs to put IoT sensors on water pumps in rural Africa. This enables the NGOs that install the pumps to track the pumps' functionality and maintain them more efficiently and in a cost-effective manner.

In a Rwanda study, only 56% of the water pumps were working consistently. After adding the SweetSense technology to track the pumps' function via cellular IoT systems and analytics, the water pumps were able to be repaired more quickly, and 91% of the pumps could be kept working on a regular basis. With projects like this from SweetSense, connected devices can help provide clean water to more people on more days, thereby improving health and well-being.

Hello Tractor helps Sub-Saharan African farmers with food production. The company works closely with its partners to create an entire ecosystem, with a sharing platform for income-generating products (tractor leasing) and affordable service offerings. This enabled more farmers to receive the services or equipment they needed to succeed. With more than 500 tractors in operation, another crucial step to Hello Tractor's success was developing a pay-as-you-go plan that farmers could afford and one that the banks and insurers could accept.

The process deployed by Hello Tractor produced greater efficiencies, higher crop yields, and a proven business model that can be implemented around the globe. Armed with this greater access to business-critical data, Hello Tractor now has plans to expand projects into Bangladesh and South Africa.





Connected devices can help provide clean water to more people on more days, thereby improving health and well-being.

AN INTRO TO THE GUIDE TO IoT FOR BUSINESS

In this book, we will focus on how the IoT ecosystem can be used by businesses. In addition to providing real-time information from devices in the field, IoT works in the other direction too (i.e., it lets companies control devices from a central location).

This ability provides everything from marketing intelligence to improved preventative maintenance. Companies can use IoT for applications as diverse as helping medical professionals care for more patients at the same time or giving retailers the ability to customize advertising to a single individual.


To get started with IoT for your business, you will need a basic understanding of what makes it all work. You don't need to be an engineer or a data scientist, but it is useful to have a grounding in the concepts of how IoT systems are connected, how they communicate, how the data is analyzed, and how this can positively impact your organization. We will present an overview on connectivity and data collection, as well as an in-depth, detailed description of the Internet of Things.

To do this, we will cover these broad topics:

- The technology that connects the Internet of Things.
- How wireless devices are networked and locate themselves.
- Different types of sensors, how they work, and what they do.
- An overview of security technologies used to protect IoT data.
- How to scale up an IoT project to immense proportions.
- Using Big Data analytics to gain insight from the IoT ecosystem.
- IoT applications and their relationship to the IoT value chain.
- Advice for managing the lifecycle of an IoT deployment.
- A view into the future of the Internet of Things.



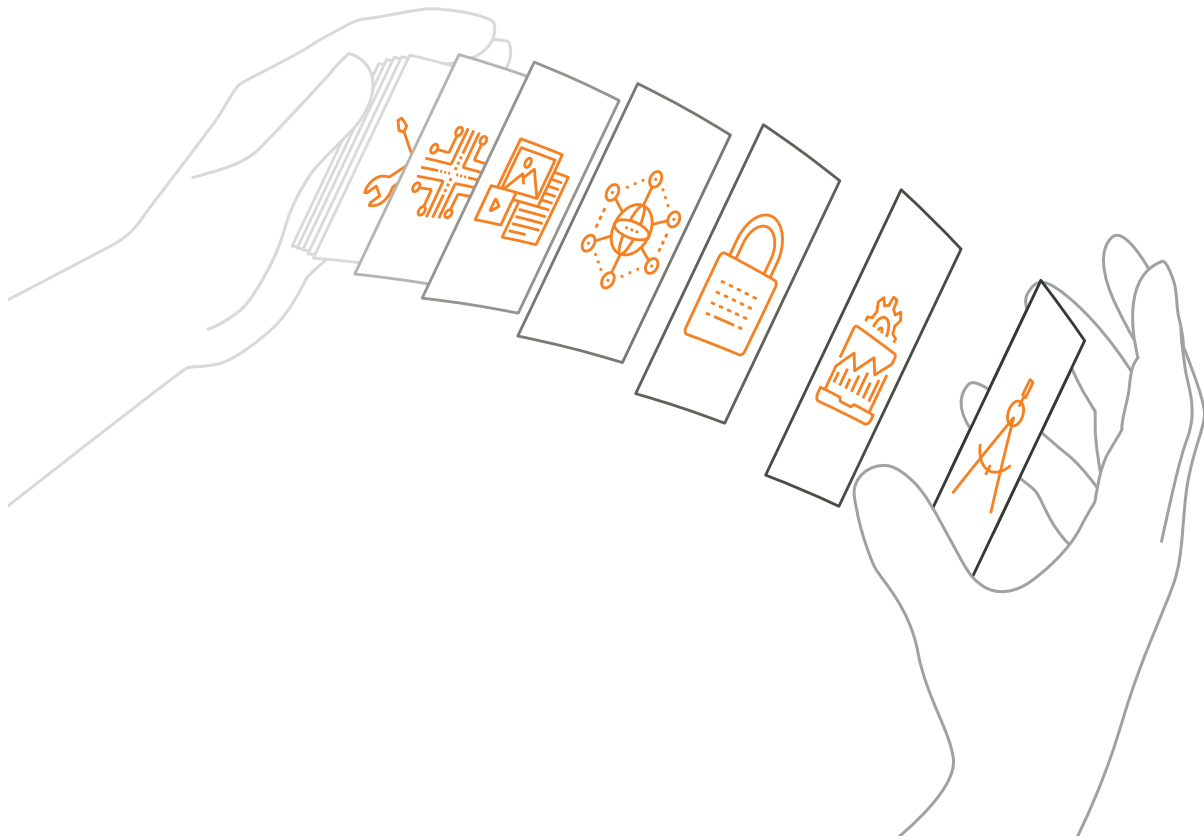
All of these IoT aspects will be addressed from an enterprise point of view for those running small to large businesses. In this guide, Aeris will look behind the scenes into how these IoT devices are run and managed, where the data they collect goes, and how it's used. If you're in the business of IoT or looking to start up a deployment, this guide is for you.



THE FUTURE OF PLATFORMS

- 15 IoT PLATFORM LAYERS AND SERVICES
- 19 FROM MANY TO ONE
- 21 ADVANTAGES OF A SINGLE VS MULTI-PLATFORM DEPLOYMENT
- 22 PLATFORM REQUIREMENTS





THE FUTURE OF PLATFORMS

Often, the ability of a company to deploy an IoT application is limited by the expertise and capability resources available to it. Engineering talent is difficult to find, and a complete IoT solution may require significant software and systems development. Experience has shown that this can lead to long schedules for deployment, which can be problematic for a variety of reasons: the company may miss the product target window, the project may be more expensive than expected, management may lose confidence in the IoT application, or future support and maintenance may tie up resources and make the IoT application fail to achieve its objectives.

In IoT, to get from the device sensors to the networks gathering data, you need an IoT platform. These days, everybody claims to have the best, the fastest, the most secure, the most adaptable platform.

IoT platforms combine many of the tools needed to manage a deployment—from device management to data prediction and insights—into one service.

Thus, it often is best to use a platform where much of the development and operational work is done by a supplier who has the necessary expertise and available capability and capacity to support the company. That supplier provides what is called a “platform” that allows the customer to focus on the objectives and development of the application rather than the mechanics of an implementation. Some platforms also may provide core capabilities of transport networks, data storage, and analytics functions that can be used more rapidly than building in-house. The supplier can support and maintain the platform in the future, including changes to support new services and new technologies.

To figure out the best solution for your deployment, look at both your company’s immediate needs, as well as for long-term deployment requirements.

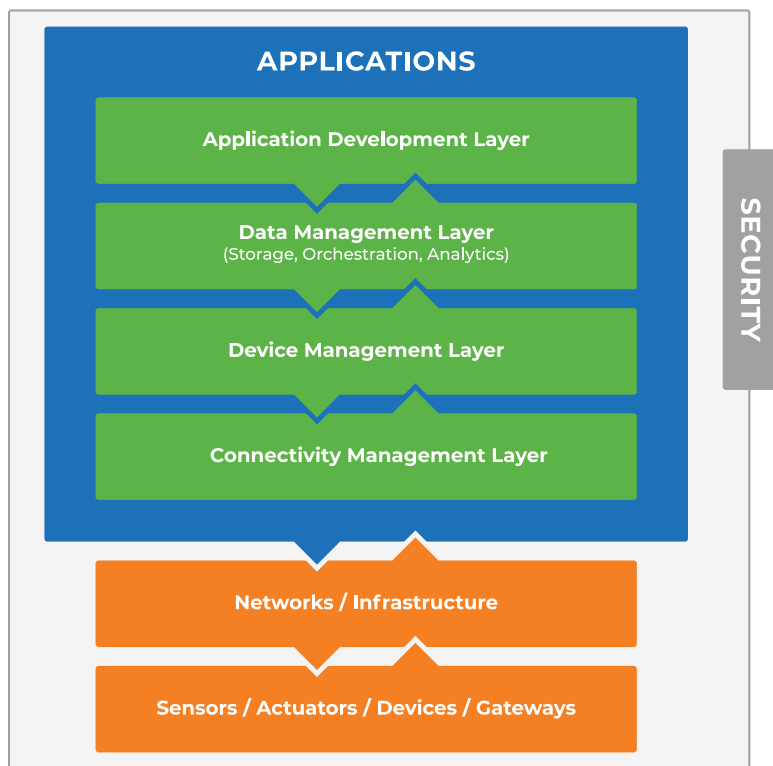
Certain criteria will differentiate platforms. These include scalability, ease of use, security, cost, third-party software integration, and a slew of additional functionality. So, what is best for your business? Well, it depends. Let’s break down some of the criteria you need to look for in a highly functional platform.

IoT PLATFORM LAYERS AND SERVICES

There are four basic layers to an IoT platform: Connectivity Management, Device Management, Data Management and Analytics, and Application Development. It is important to look at all four layers, along with an overarching layer of Security.

Connectivity Management

This layer of an IoT platform includes the connectivity types supported, the network protocols that are available to devices, the coverage and footprint of the connectivity (whether it is regional, national, or global), the device provisioning and subscription management, and rate, usage, and billing management.



Beecham Research, 2017

Device Management

The IoT Device Management platform can offer elements for automatic discovery of new devices and their configuration, device health monitoring and updates, library support for a diverse set of devices, scalability to allow growth—beyond the 100,000 unit mark without problems, and edge analytics where the platform processes data and potentially takes action at the device and / or gateways.

Data Management and Analytics

The Data Management and Analytics platform can offer data storage capabilities (including cloud services and data centers), orchestration for disparate sources of data, analysis and visualization tools, and advanced analytics (such as predictive analytics).

Application Development

The Application Development platform may offer tools for the development and deployment of applications, including dashboards and portals, application programming interfaces (API) to the dashboard, data and cloud computing services for critical functions, and modeling and simulation tools to enable application development without actual deployment.

Security

Any Security platform that is expected to work with an IoT application must provide services for access, authentication, and authorization of a device (or groups of devices) for the user (or multiple users), content protection via data encryption, and key management, as well as gateway security protection and management (including updates).



FROM MANY TO ONE

A platform needs to provide a single interface to manage and monitor usage across all deployed devices, regardless of the underlying carrier. This visibility and control can lower operational costs of IoT solutions. Combining real-time information with actionable device controls provides command throughout the entire deployment lifecycle.



Manage device status—Provision, activate, suspend / unsuspend, and cancel service on devices.



Monitor usage in real time—Run on-demand reports for network traffic, billing information, and network registration events.



Set customizable alarms—Configure real-time alerts and alarms on device activity for customized proactive control.



Understand device behavior—View daily and monthly summaries with detailed reports and rated billing events.

Think of one platform to manage multiple technologies, with additional capabilities to identify problems faster, resolve issues quicker, and control costs at every stage of the device usage.

Businesses from every sector are requiring single-platform functionality where all technologies, as well as connectivity from multiple carriers, could be implemented, viewed, and managed. The platform needs to provide API integration for visibility and management processes; seamless third-party integration; a simplified process that would reduce costs; and an overall reduction in operational complexities.

Using a single platform to manage the applications provides distinct advantages from an on-going operational cost and complexity perspective. Overall, the differences between single- and multi-platform deployments can be quite extensive, with single platform deployments being much more conducive to success. Devices may be deployed globally in multiple networks (for example, cellular and non-cellular technologies) and managing the entire deployment, including devices, the data, and analytics, particularly at large scale, is best done with a “single pane” view of the solution. (See table on the following page for even more details.)

The bottom line—
IoT needs security
at the design
stage, and not as
an afterthought.

Today's advanced IoT platforms include a cloud-based software solution for connecting to and managing the connected features of IoT devices. Built using a micro-services architecture, the platform should include an extensive set of software services, packaged as micro-containers, which are functional building blocks that can be combined in a multitude of configurations to create application-specific variations. The power of the platform comes from its customizable ability to address multi-variant complexity.

A global platform needs to allow for regional market variations (brand, location, and more), thereby redefining and expanding the ability to gather, analyze, and react to volumes of data. This concept could apply to multiple business sectors (automotive and healthcare will lead the way, but other sectors will engage this technology in the future).

ADVANTAGES OF A SINGLE vs MULTI-PLATFORM DEPLOYMENT

	Multi-Platform Deployment	Single-Platform Deployment
Scalability	Deployment becomes a collection of individually managed groups	Single management interface worldwide simplifies large scale operations
Complexity	Operations teams need to be trained on and manage multiple systems	One portal, one set of operational processes to identify and resolve issues
Efficiency	Operations teams take longer to identify issues requiring larger teams at a higher cost	Issues are identified faster and resolved quicker, by fewer resources
Visibility	Lack of consistent visibility into devices and usage patterns impacting end-user experience and billing	Holistic reporting and analytics provide complete insight into entire deployment
Support	Multiple support processes to follow with no standardized SLAs	One support process to follow for all issues and all devices

PLATFORM REQUIREMENTS



Global Scalability—Initially, global scaling must take into account regional languages, local regulations, and differing access capabilities. Make sure the documentation is available in multiple languages. Likewise, partner with a local company versed in the target country’s customs, language, and laws. This will help resolve problems quickly and efficiently, saving both time and money.

Since the nature of IoT is to deal with millions of devices, the upgrading capacity becomes critical in pushing new services to consumers in as little time as possible. Seamless and timely upgrades must be part of any long-term solution, with over-the-air (OTA) updates to limit the need for human intervention (thus materially reducing costs). Additionally, in a global deployment, having one portal (“single-pane-of-glass” view) to configure and manipulate all devices is key to offering up-to-date and continuous device management and configuration.

Also, understand device certification and equipment requirements for each country. Just because a device is certified in one country does not mean it will pass certification elsewhere. The bottom line—find a service provider with international coverage and roaming capabilities, local language support, 24x7 device monitoring and management, and the ability to administer devices from any location.



Platform Security—Securing an IoT platform requires an end-to-end approach, from physical devices and sensors, to data connections, to host systems, to the services and data stored in the cloud. And while security risks can never be completely eliminated, find a provider with the tools and expertise to mitigate these risks with the responsible development of IoT applications.

The bottom line—IoT needs security at the design stage, and not as an afterthought.



Service Flexibility—IoT solutions need to address a key requirement of every global business: regional flexibility. A platform must allow deployment of a global program across a multitude of public and private clouds in a manner that maximizes operational efficiency and maintains local flexibility and autonomy while also ensuring compliance with local regulations.



Customer Experience—When it comes to IoT solutions, customers require and expect on-going, long-term reliability without the need for constant human touch. Support needs to be agile and responsive to customer concerns and needs. And the platform itself needs to make IoT deployment management as simple as possible, thereby saving work-hours and resources.


Modern platforms address these issues with a services delivery architecture that optimizes an end-to-end system—from wireless connection to telematics applications—while simultaneously addressing key concerns around hardware power management, cost structures, and visibility.

Today, there are hundreds of companies, as well as numerous startups, concentrating on IoT platform development.



Your IoT platform has to ensure the largest possible geo-coverage of countries and continents, along with the ability to develop services and enhance products seamlessly. The platform has to provide a superior combination of coverage and operational time, while offering a heightened level of security.

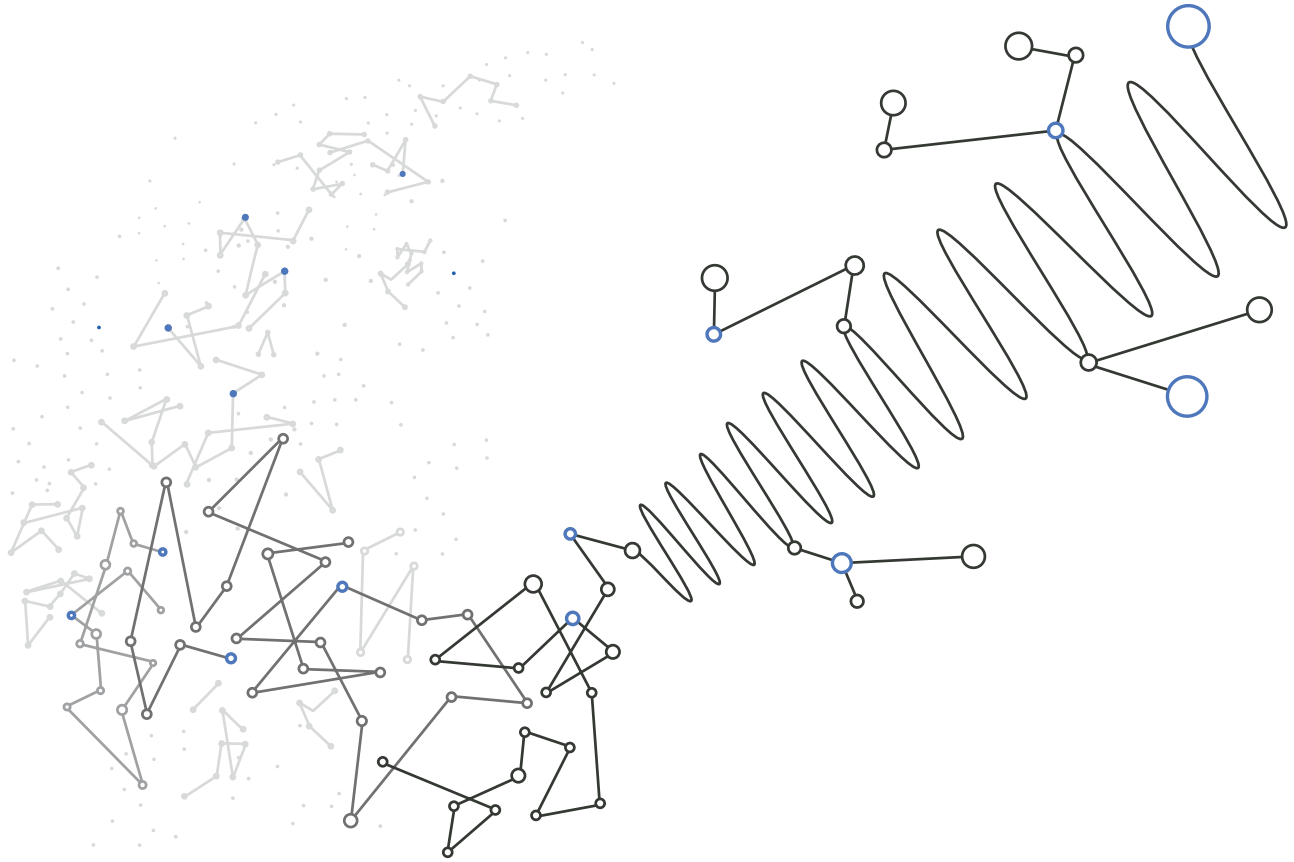
Seek a solution provider that has an end-to-end modular technology platform that enables customers to make the pivotal change from unconnected product to connected services offerings. This journey is a challenging and complex process and you need a vendor that has the knowledge, experience, and technology to implement, accelerate, and optimize these services, which include connectivity, storage, analytics, and service layer functionality.



IoT CONNECTIVITY: TYPES AND CHOICES

27	BASIC INTERNET CONCEPTS
28	CHOICE OF CONNECTIVITY
32	PROPRIETARY PROTOCOLS
34	STANDARDIZED PROTOCOLS
37	TYPES OF CELLULAR TECHNOLOGIES
44	CELLULAR FALLBACK
45	HOW TO DETERMINE LOCATION
47	GLOBAL POSITIONING SYSTEM





IoT NETWORK TECHNOLOGY

To understand how the Internet of Things communications work, you need a basic overview of the technology used for the internet. While technology always is evolving, certain principles are common to how networking functions. What changes more frequently are the tools and protocols used to access the network, such as modems, cellular radios, transmitters, and more.

BASIC INTERNET CONCEPTS



IP—Traffic on the internet uses the Internet Protocol (IP) to transmit data. This communications protocol has a routing function designed for internet connectivity. IP is used to route data packets across an IP network from a source host to a destination IP address. Every node in such a network has an IP address, a unique numerical label. The computers and printers in your office generally have private, local-area network IP addresses, while websites, such as Aeris.com, have public IP addresses.



Packet—Data travels across an IP network in packets. Each packet has both a source and destination IP address, but many packets may be needed to make up one complete “item”. For example, a single email message can be comprised of many different IP packets that, when assembled by the remote network for the recipient’s email program, make a complete piece of mail. A webpage retrieved by your browser also is comprised of multiple packets.



Router—A router connects one network to another. For example, your home or office wireless router connects the internal IP network in your home or office to the public internet via an Internet Service Provider (ISP). Your ISP connects to other providers and internet backbones using routers.



Modem—A modem is a shortened term for “modulator demodulator”. The modem modulates signals to encode digital information and demodulates the received signal to retrieve the information. Wireless broadband modems are a popular way for smartphone and laptop users to get internet connections. Early wireless modems used the 2G and 3G cellular standards, but most have moved to the faster 4G LTE technology, which rapidly is becoming available around the world.



Speed—Internet speed is measured in megabits per second (Mbps). For example, Netflix HD video typically requires five megabits per second for good video quality viewing, although its service will work at speeds as slow as 0.5 Mbps.



The top countries for the fastest average internet connection speeds are led by South Korea (28.6 Mbps); Norway (23.5 Mbps); Sweden (22.5 Mbps); and Hong Kong (21.9 Mbps). The U.S. enters the list in tenth spot with averages speeds of 18.7 Mbps.¹

¹ Akamai Q1 2017 global average connection speeds ranking.

CHOICE OF CONNECTIVITY

The internet can be accessed in many ways, depending on your device and application. There are pros and cons to each form of connectivity technology, particularly when implementing a large IoT project.

Internet Service Providers

An ISP connects offices and homes to the internet by taking their network traffic and forwarding it to other networks until it gets to the desired destination. An ISP could be, for example, Telstra in Australia. But it doesn't stop there because an ISP has to connect to other ISP networks. For example, while Telstra runs a large internet network in Australia, it still has to connect to other networks within the country and around the world.

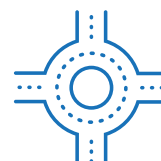
Tier 1, Tier 2, and Tier 3 networks form the internet's backbone. ISPs, such as Telstra, connect to those networks, which become the principle routes for internet data transmissions around the world.

Wireless operators, such as Aeris, connect IoT deployments to the internet or private networks in a similar fashion. A wireless operator has a cellular network that uses fixed base station radios at cell towers instead of wires to transmit signals from the cellular devices into the network. Much like ISPs using other ISPs, wireless operators also can connect to Tier 1, Tier 2, or Tier 3 networks. This is how they deliver traffic on the wireless network when a mobile device requests data.

Wired and Wireless IoT Connections

A home, office, or IoT-networked device can be connected to the internet either via a wired or wireless connection. If the connection is wired, it generally is connected directly into an internet router, and the device needs to remain stationary. A device with a wireless connection can have a cellular modem, a Wi-Fi router, or other connectivity technology, which, among other things, lets the device be physically mobile.

Wired connectivity was common in the early days of M2M systems. For example, many factories installed wired systems for supervisory control and data acquisition. For business and residential security systems, alarm panels could use telephone circuits to communicate events—like a burglary or fire—to central monitoring stations.



Connectivity, however, depended on where the ISP's lines could extend, with setups possibly being complicated. These early applications tended to be purpose-built, meaning each industry and company developed its own devices and software systems from scratch for a specific purpose.

The 1990s saw a move towards using wireless radio technologies in these applications. Ademco Corporation (now a division of Honeywell), a leader in intrusion and fire detection systems, began to build out a private radio network to address this need. In 1995, Siemens introduced the first cellular radio module for data transmission applications. Very shortly afterwards, Aeris introduced its MicroBurst™ data services using the control channels of the Advanced Mobile Phone System (AMPS) cellular service. Ademco became the first major customer to deploy M2M devices using this transport.

These new technologies enabled machines to be free from wires, and more IoT functions were possible in different industries, including consumer products. For example, in 1995, OnStar® became one of the first connected car systems, offering a mix of safety services and entertainment options. Fleet and container tracking solutions similarly made use of mobile telematics for the trucking and railroad transportation industries. In addition to being mobile, cellular connectivity could extend application reach to more remote locations than wired networks could allow.

By the 2000s, changes in cellular technology introduced digital cellular networks with features such as Short Message Service (SMS), General Packet Radio Services (GPRS), and 1 Times Radio Transmission Technology (1xRTT). However, there arose two competing types of digital cellular, CDMA and GSM, with different industries choosing differing solutions. In the U.S., for instance, the automotive and trucking industries mostly chose CDMA devices, while the alarm and security industries generally picked GSM.

Looking Forward

The future holds promise for more varieties of wireless data technologies, including wider adoption of 4G LTE, and, eventually, 5G in the next few years. Short-range data transport methods, such as Bluetooth, ZigBee, and 6LoWPAN, may be used to augment long-range cellular in some applications. We also are seeing the commercial deployment of Low Power Wide Area Networks (LPWAN) that provide long-range communication similar to traditional cellular, but consume much less power.

New technologies enabled machines to be free from wires.

The future holds promise for more varieties of wireless data technologies, including wider adoption of 4G LTE, and, eventually, 5G in the next few years.

Low Power Wide Area (LPWA) Networks

Several new LPWA connectivity technologies, cellular and non-cellular, licensed and unlicensed, are trying to win the hearts, minds, and pocketbooks of companies worldwide. Overall, these connectivity strategies will afford a diverse range of business sectors to seek out operational efficiencies and competitive advantages through collecting, storing, and analyzing business-critical data at levels of granularity previously unseen.

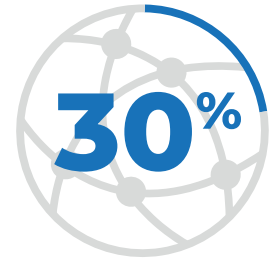
This shift to LPWA is promising to bring in many more industries under the IoT umbrella. For many industrial functions that require sensors or devices obtaining and sending only small amounts of data, LPWA offers the best low-cost option. Additional factors are weighing heavy in the promotion of LPWA networks, including less expensive devices that are reaching the market; low throughput for long- or short-distance transmission; data transfers that support small blocks of data intermittently sent; and the ability for extended coverage, both indoors and underground. What this all means is extended reach, a much lower cost of entry, and a much lower total cost of ownership.

For businesses, however, one size, or sometimes, one technology or one cost structure, doesn't fit all. So, the growth of multiple LPWA IoT solutions is a good thing.

How LPWA Networks Work

LPWA networks are designed for IoT and M2M applications that have low data transmission rates, need long battery lives, can provide low-cost services, sometimes operate in remote or hard to reach locations (underground or geographically dispersed), and be easy to deploy across basically every business sector, including manufacturing, automotive, energy, utilities, agriculture, healthcare, wearables (for humans or animals), or transport.

Present-day cellular mobile technologies are designed to work on costlier consumer-oriented networks where a premium is placed on fast connections that can transport large amounts of data. Low-cost LPWA networks, however, can support devices requiring low mobility, low-power consumption, long-range abilities, and heightened security. One of the benefits of LPWA is that data transfer rates, as well as power consumption, are very low. Device connectivity in this case requires less bandwidth than standard cellular, which means that LPWA networks can operate with far greater power efficiency.



It is predicted that by 2019, 30% of connected devices will be on LPWA networks.

Additionally, LPWA networks can support more devices, at a lower cost, over a larger coverage area than consumer mobile technologies.

Initially, IoT services relied on licensed cellular, wireline, and satellite networks for wide area connectivity requirements. These, however, were not a good fit for widespread IoT usage due to excessive power consumption and complex protocols that lowered battery life. Recently, to help alleviate these issues, several (more) LPWA alternatives have appeared on the market. These networks, generally, are more business friendly, with low data rates, extended battery life, and extended coverage.

Technologies, Protocols, and Players

There literally are dozens of participants in the burgeoning LPWA sector. Basically, they can be broken into two overarching categories: standardized and proprietary. The differences between the two categories are fairly basic. Standardized LPWA connectivity runs in a licensed spectrum, generally compatible with existing cellular standards. Proprietary technologies, run in an unlicensed spectrum. They can get to market faster, but adoption of this technology still is in question. Below, we will define some of the technologies at play in the IoT sector, along with some primary advocates of each protocol.



PROPRIETARY PROTOCOLS

Random Phase Multiple Access (RPMA)

RPMA is a low-power, wide-area channel access method used for IoT and M2M communications. RPMA employs direct sequence spread spectrum (DSSS) modulation to access the best signal for both the network and its devices. It is IEEE 802.15.4k compliant; uses a globally available, cost-free unlicensed spectrum; requires low-power support, thereby extending battery life; and provides high network capacity. Additional selling points of RPMA include extended coverage with high capacity for multi-million-node networks. RPMA uses standards-approved algorithms for both device and messaging security.

In the RPMA sector, companies such as Ingenu (formerly OnRamp Wireless) are using RPMA protocols designed specifically for wireless IoT communications. The company's commercial deployments are in the unlicensed 2400 MHz range. This service initially was deployed for private networks but now is publicly available in select areas of the U.S. RPMA has excellent noise immunity for range and throughput, with a high link budget (less interference). Data is highly secure (using AES 128 bit encryption) and the product has seen some success in utility markets with private networks.

On the downside, this technology is playing catch-up in the U.S., has radio costs that, generally, are higher than other protocols, and, as a proprietary technology, deployment and support are limited.

Ultra-Narrow Band (UNB)

UNB technology transmits over a low bandwidth, in a very narrow license-exempt radio spectrum channel (less than 1 kHz) to achieve long-distance data links between a transmitter and a receiver. UNB is fully bidirectional, meets the long-range, low-cost needs of business connectivity, does not rely on other networks, and already is in use in multiple smart applications (lighting, meters, etc.) UNB is gaining success by combining long-range connectivity with an extended battery life (up to a decade).

As a major proponent of UNB, Sigfox provides a software-based LPWA communications solution, where all the network and computing data is managed in the cloud. Working in the unlicensed spectrum, its UNB proprietary solution uses a simple protocol with slow data speed. Originally France and European Union focused, Sigfox started to deploy in the U.S. and elsewhere in early 2017 and has partnered with a number of firms for advancing UNB technologies, including Texas Instruments, Silicon Labs, and ON Semiconductor.

UNB provides end users with low device costs and low energy consumption. It employs a simple API to integrate radio modules, and this protocol has several chip suppliers. As stated, though, this technology is not fully engaged in the U.S. Its message size of 12 bytes (by design) limits its applications, and it has shown some reliability issues.

LoRaWAN

The LPWA network specification LoRaWAN focuses on secure bidirectional communication, mobility, and localization services and is a protocol supported by the LoRa Alliance. It is intended for wireless, battery operated devices. This open standard, developed by Semtech, and supported by IBM, SoftBank, and several other carriers, operates in unlicensed spectrum, using narrow spectrum in the 868-915 MHz ISM band, up to 500 kHz bandwidths. LoRaWAN uses an adaptive data rate (along with radio frequency output) for each device so as to extend battery life and increase overall network capacity.

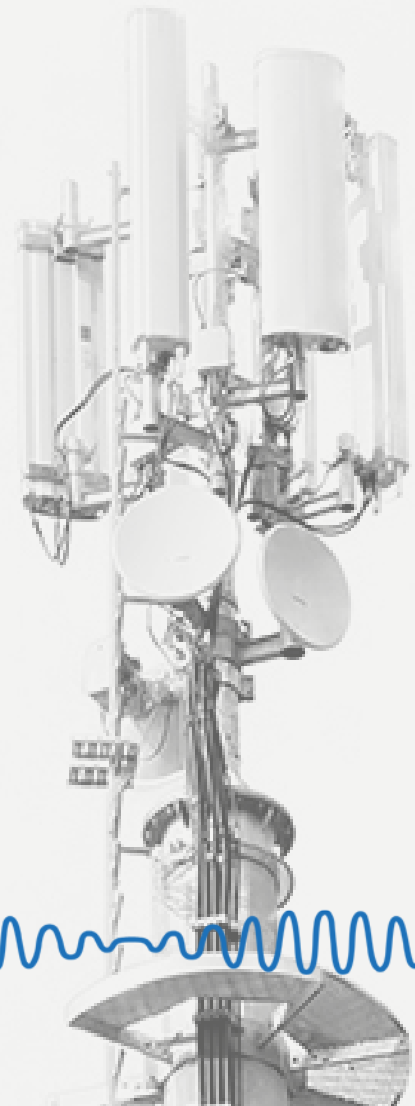
LoRaWAN data rates range from 0.3 kbps to 50 kbps. This protocol uses spread spectrum technology, along with virtual channels, to separate cross-channel interference. Additionally, LoRaWAN provides encrypted data for higher security levels, gateways and radios from multiple suppliers. However, its message limit, defined by the user, can be no more than five seconds in length to meet compliance requirements. Its radio costs still are high and, as of the publication of this book, U.S. coverage is limited to select areas of the U.S.

The LoRa Alliance is lobbying hard for its global protocol, LoRaWAN, in LPWA networks. IBM is using LoRa's wireless sensor network with its Long Range Signaling and Control (LSRC) software, as well as its IoT cloud-hosted service, to create large-scale IoT deployments. Its open standard is supported by several organizations and businesses, and IBM's global influence could be the thumb on the scale, providing the business recognition, along with a low price point, for this solution's continued growth. Efforts to deploy American public networks are underway, but coverage still is limited today.

Nwave

Nwave uses ultra-narrow band (UNB) radio technology, combined with advanced software defined radio (SDR) techniques, to provide a communications network for IoT. Ultra-narrow band, which operates in internationally available and unlicensed sub-1 GHz radio spectrum (ISM bands), allows for strong signal dissemination, giving enhanced in-building penetration and range, while using minimal power.

Unlike some other IoT communications technologies that require the use of a mesh network, UNB also is highly scalable, allowing for high capacity networks with a simple star architecture, whereby devices communicate directly, and securely, with base station transceivers.



STANDARDIZED PROTOCOLS

IoT industry standards, as set by the 3rd Generation Partnership Project (3GPP), recently saw more protocols join the LPWA sector. These include LTE-M (also known as LTE-CatM1 or LTE-MTC) and Narrowband IoT (NB-IoT). Each wants the crown of global LPWA standards coming their way. Each has an argument for dominance.

LTE-M

LTE-M is a bidirectional, standards-based protocol within a dedicated spectrum. It provides carrier-grade security, long battery life, low data needs, and low-cost modules. This protocol has many active followers, including Altair, Ericsson, Qualcomm Technologies, WNC, and Xirgo, as well as a host of U.S. and international carriers.

One of the strengths of LTE-M is that it does not need new infrastructure as it can piggyback on existing LTE networks. What that means is that a carrier can update software on its network and get LTE-M functional. LTE-M, however, is a much simpler product, only using 1.4 MHz of the channels instead of 20 MHz.

Additionally, using an extended discontinuous repletion cycle (eDRX), the data collection devices can transmit data on a non-continuous schedule, as set by the end user. The device, when not sending data, is not off, but just asleep. When data is scheduled to be sent, the device does not need to be re-activated to join the network, it just wakes up. Having intermittent data send-schedules, which are not active 24x7, can save battery life, leading to significant cost savings. Data rates for LTE-M are somewhat higher than NB-IoT, but it can transmit larger blocks of data. LTE-M had an American roll-out in 2017.

Extended Coverage GSM (EC-GSM)

Extended coverage GSM is a standards-based, LPWA technology. It is based on eGPRS and designed as a high-capacity, long-range, low-energy, and low-complexity cellular system for IoT communications. The optimizations in EC-GSM from existing GSM networks can be accomplished via a software upgrade, ensuring coverage and accelerated time to-market. Much longer battery life (up to 10 years) will enhance usage in multiple use cases.

The first commercial launch was in 2017. Supported by all major mobile equipment, chip set, and module manufacturers, EC-GSM networks will co-exist with cellular mobile networks, along with all the security and privacy mobile network features, including user ID confidentiality, entity authentication, and data integrity.

Narrow Band IoT (NB-IoT)

NB-IoT is the newest entry to the IoT scene. With its standards-based LPWA technology, NB-IoT has a global reach with better bidirectional data than any of its unlicensed competitors. And, unlike LTE-M, NB-IoT is based on Direct Sequence Spread Spectrum (DSSS) modulation, which 'spreads' the signal so as to reduce interference. It also might make it a bit harder to go national (since it can't yet hook into a typical LTE network).

NB-IoT has several large organizations, including Huawei, Ericsson, Qualcomm, and Vodafone, actively involved with this standard. Additionally, NB-IoT is supported by all major mobile equipment manufacturers and can work with 2G, 3G, and 4G mobile networks, so it enjoys the heightened security of mobile networks, including user ID, authentication, data integrity, and more.

The NB-IoT LPWA solution is optimized for applications that need to communicate small amounts of data over long periods of time. NB-IoT results in lower latency with a higher transmit power limit (200 kHz bandwidth), which improves range and reliability, even underground or inside buildings. And since it operates in a

licensed spectrum, it is secure with highly reliable data transmission, assuring a high quality of service.

NB-IoT devices and hardware are at the lower end of the cost spectrum and improved efficiency helps batteries last more than a decade, allowing for long-term IoT application deployments. With its simpler underlying technology, costs for NB-IoT modules will continue to decrease as demand increases. The technology roll-out, with a commercial module and network, started in 2018 and will continue through 2019. In the U.S., the present expectations are that all of the big four wireless carriers will deploy nationally, with Verizon and T-Mobile hitting the market first, followed by Sprint and AT&T. Globally, a large number of carriers have begun NB-IoT deployments in their coverage footprints to support IoT applications.



Address Spaces / Numbers

Too Many Internet Devices for IPV4

Due to the explosion in the number of websites, mobile devices, and always-on IP connections (the latter of which is crucial to future IoT deployments), the internet's governing bodies realized that the IPv4 address space would not be sufficient over the long term.

Luckily, the shortage noted in 2011 has not had a serious impact because of techniques such as Network Address Translation (NAT). This allows a router to share the same external public IP address, or set of public addresses, for all the traffic generated by systems on the internal network. Because of NAT, many internal systems can share a common IP address for external internet access.

But the long-term solution for accommodating the billions of devices constantly being added to the internet, especially with IoT applications, is to upgrade the IP address space to a much larger number range.

The World Is Moving to IPV6

The problem of not having enough IPv4 address numbers will be resolved when the internet world moves to IPv6, where the total address space has been expanded to 128 bits (from the 32 bits used in IPv4). This allows 2 to the power 128 (or approximately 3.4×10 to the power 38) IPv6 addresses.

Although not yet fully deployed across the internet, IPv6 networks already are in use by many large corporations and websites. For example, Google and Facebook have provided access to their systems in IPv6 networks.

Ultimately, every device and router will use IPv6 addresses to access the public internet. In the interim, gateway systems provide address translation functions, thereby allowing older IPv4 systems to access future IPv6 networks.

TYPES OF CELLULAR TECHNOLOGIES

This section provides an overview of the cellular technologies available to IoT devices and applications for long-range data transmissions. These cellular technologies are evolving and will continue to change over time. You should assume that new cellular technologies completely will replace existing deployed technologies over time, so plan the device and application lifecycles accordingly.

Brief History of Cellular

Cellular service has evolved over time. Often, a fairly major change in the technology would render a previous technology incompatible and necessitate a replacement of the radios and handset, along with changes in the network to support the new radios.

In the cellular industry, these major changes loosely are termed “generations” to distinguish and summarize their technology, the protocols used, the network changes, and the commercial deployment phases.

Analog Cellular

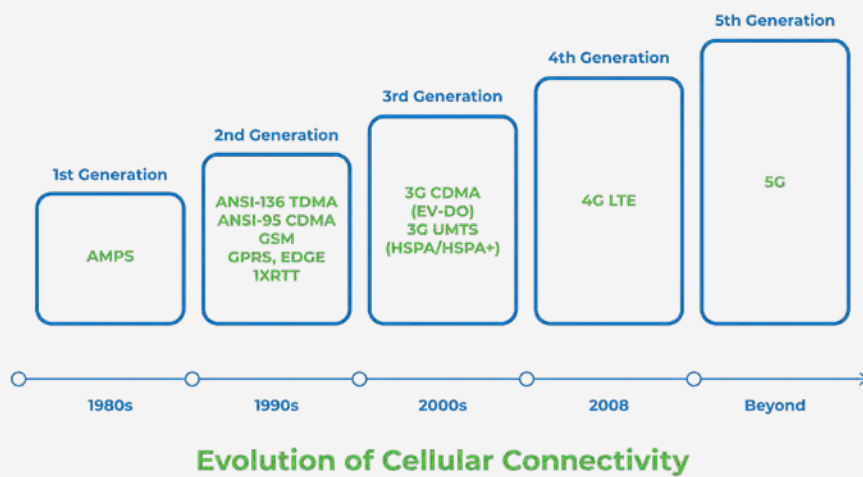
The first cellular service was an analog cellular system, later termed First Generation (1G). In North and South America, this was the Advanced Mobile Phone System (AMPS). It was deployed in the U.S. in the early 1980s and was shut down in February 2008.

AMPS used radio frequencies (spectrum) distinct from other wireless services. In particular, the technology used relatively low-power transmissions, which restricted the distance of the radio signals, to reach a tower (also called a base station) where the voice call could be sent into the landline telephone system.

New cellular technologies completely will replace existing deployed technologies over time so plan the device and application lifecycles accordingly.

This allowed re-use of the radio channels beyond a particular distance from a tower—each tower received and transmitted only to the cellular radio devices within that range. Grouped into cells (hence, the term “cellular”) resembling a beehive, the tower radio did not communicate with devices outside its cell. Cellular devices communicating in remote cells could use the same radio channels (i.e., re-use the frequencies) without interfering with calls in the closer cell.

Eventually, AMPS and other analog cellular services were discontinued in most parts of the world (in the U.S., this was the “AMPS Sunset” in February 2008).



TDMA and GSM

To maintain backwards compatibility with AMPS in the early deployments, technologists in the U.S. used a mechanism to slice each AMPS radio channel in time, hence the general term for the protocol: Time Domain Multiple Access (TDMA). Humans are unaware of the missing “times” when the channel is used for other voice calls, as long as the duration of the missing time is short enough. The TDMA protocol is quite successful at this function.

The standard deployment was called EIA-136 TDMA (eventually ANSI-136 TDMA), and it improved the efficiency of the channel by a factor of three (since each call only used the channel one-third of the time). Essentially, each channel now could support three TDMA voice calls simultaneously rather than one AMPS voice call.

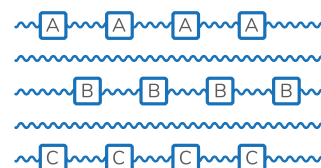
In Europe (and eventually most of the world), another TDMA approach was used, termed the Global System for Mobile Communications (GSM). The European and U.S. TDMA technologies were incompatible due to spectrum allocations and channel differences, which meant that a GSM cellphone could not operate in an ANSI-136 TDMA network and vice versa.

GSM rapidly became popular in Europe and in other parts of the world and, with a far larger deployed base of cellphones, the economies of scale meant that GSM cellphones rapidly became lower in cost than ANSI-136 TDMA cellphones. Thus, the operators in North and South America eventually abandoned ANSI-136 TDMA in favor of GSM to take advantage of this reduced cost.

CDMA

In the 1990s, another new digital protocol was deployed, mostly in Asia and North America. Rather than using TDMA encoding, the digitized human voice bits are combined, or multiplexed, with “codes” using a mathematical algorithm. Thus, this encoding protocol is called Code Division Multiple Access (CDMA).

The combination of voice bits combined with codes allows the data to be transmitted over a single, wider channel. The bits essentially are “spread” across the spectrum width of that channel, and it is thus a “spread-spectrum” communications system. Over the long term, however, CDMA struggled with adoption in other parts of the world, and this led to a much smaller prevalence of CDMA devices that remains today.



Data Transmissions

When cellular systems became digitally encoded, it was natural to consider treating the transmitted digital bits as something other than human voice-encoded bits. This allowed the deployment of data transmission services for purposes other than human voice. This included communications from mobile radio devices (data handsets) and data cards for mobile computers (laptops) to access the increasingly important internet and the World Wide Web.

The mechanism for treating the digital bits as application data, rather than human voice, was different in the deployed technologies.

2G GSM DATA: GPRS, EDGE

GSM introduced a practical data transmission technology called General Packet Radio Service (GPRS), followed by an improvement called Enhanced Data Rates for GSM Evolution (EDGE) with higher throughput.



These technologies were popular for cellular data communications, although the throughput rates are extremely slow compared to today's smartphone needs. In IoT applications, however, where the throughput requirements were lower, GPRS has been a sufficient technology for low-data rate transmissions. Thus, GPRS was widely used around the world for early cellular IoT and M2M applications.

2G CDMA DATA: 1XRTT

Like GPRS in GSM, CDMA operators in many countries deployed a data transmission technology called 1x Real Time Transmission (1xRTT). This was faster than GPRS in its base throughput rate and provided a reliable, extensive coverage data network for IoT applications.

In the U.S., the wide availability of 1xRTT made it an easy choice for physically mobile applications, such as the automotive and trucking industry, that needed coverage across the continent. The early deployment and expansion of CDMA and 1xRTT led to excellent cross-country coverage.

However, the complexity of the CDMA data encoding protocol compared to TDMA resulted in a higher cost for the radio modules, since chipsets for CDMA radios are more complex. Although CDMA technologies for data transmission continued to evolve through 3G and 4G, very little was used for IoT applications past the 2G CDMA 1XRTT iteration.

3G UMTS (HSPA / HSPA+)

Over time, it became clear that 2G GSM voice and data transports, like GPRS and EDGE that used the TDMA encoding protocol, were not spectrum efficient. The cost of adding new spectrum continued to increase, as national governments began auctioning new spectrum for smartphone data uses.

3G UMTS was developed to solve these and other issues. However, in most IoT / M2M applications, the performance and throughput of the 3G technology outpaced the needs of the M2M marketplace at the time. The rapid development of 4G LTE networks outpaced the expansion of M2M throughput needs, and, therefore, many 2GSM IoT / M2M applications over the last several years leapfrogged 3G technologies and went directly to 4G LTE. The final choice generally becomes a function of the cost of available radio modules and service coverage.

3G CDMA (EV-DO)

On the CDMA side, 3G data standards were improved substantially to enhance their data throughput rates with EV-DO Rev. A and EV-DO Rev. B. However, much like 3G UMTS, 3G EV-DO has not been used extensively for IoT / M2M applications for the same three key reasons: the higher throughput (compared to 1xRTT) is not strictly required, 1xRTT coverage and availability in the U.S. was excellent, and the EV-DO radio module costs are higher.



4G LTE

One of the limitations faced by 3G technologies was that they used fixed-width channels. With the ever-increasing number of smartphone data users, the availability of wireless spectrum has created many new bands that are not always optimally usable by 3G technologies. National governments have auctioned a large number of new bands for smartphone users.

To use these new bands, the standards entities developed a new technology for more flexible spectrum use. Since they also had the opportunity to select the encoding protocols to use these new bands, Long-Term Evolution (LTE) was designed to use a new protocol called Orthogonal Frequency Domain Multiple Access (OFDMA). The specific encoding details of OFDMA is beyond the scope of this book, but it has been termed a Fourth Generation (4G) technology, since it is quite different from 3G and also meets some of the original performance requirements set for new cellular implementations under the umbrella of a 4G service.

What is quite important, however, is that LTE is very flexible in terms of channel widths that can be used and, thus, the available spectrum bands can be partitioned into smaller blocks with greater ease. And it also allows existing spectrum to be partitioned into multiple blocks, which can allow an operator to deploy 4G without having to entirely remove older technologies.

This flexibility comes at a price. There are more than 40 bands available for LTE use, and countries have not auctioned or made available the full set of possible bands. Indeed, some bands may be impossible to use for LTE in certain countries because they are dedicated to other uses.

Therefore, handsets that can be used for LTE everywhere must support a number of different bands, and the addition of each band adds cost, since filters and power-amplifiers inside the radios must support each band. For IoT / M2M applications, this can increase the overall cost of the radio module substantially. Smartphones can absorb the higher cost of multiple band support, since it is a smaller percentage of the overall cost of the phone.

LTE also introduces the concept of categories (CAT) to define a set of performance metrics that are dependent on specific parameters such as the number of spatial layers, antennas, and protocols. Originally defined as CAT-1 through CAT-8, these provided a different range of performance, from 10 Mbits/sec download speeds in CAT-1 through 1200 Mbits/sec downloads in CAT-8.

Chipsets now
are available
and are
substantially
lower in cost
because of
the reduced
requirements.

Most LTE smartphones use CAT-3 and CAT-4 to provide data rates that are sufficient for power users but needs continue to grow and higher categories (through CAT-18) have been added. For most IoT applications, CAT-1 radios would provide sufficient performance, but originally were not developed since the LTE chipsets with CAT-1 support were not deemed adequate for smartphone users. However, recent developments in LTE chipsets have allowed manufacturers to release CAT-1 modules for IoT and M2M applications.

The standards bodies also defined CAT-0 radios for LTE that have reduced performance and network requirements, although CAT-0 appears to have been skipped by module manufacturers in favor of the newer CAT-M and NB-IoT technologies. Chipsets now are available and are substantially lower in cost because of the reduced requirements. A subset of service providers has started to roll out network support for these categories, and this support can be expected to grow over the next several years.

5G: Next Generation Coming Soon

Due to hit the marketplace in 2020, 5G has the potential to advance and expand the IoT industry by the addition of significant improvements and greater bandwidth accessibility.

These improvements include expanding the use of spectrum for IoT; greater download speeds; more than 1000 times the capacity of 4G; reduced latency; lower battery consumption; advanced functionality when compared to earlier protocols; and much more. But there still are many issues to resolve before acceptance is assured. For instance, 5G might not be supported in many parts of the world due to the huge speed increase; radio signal issues; compatibility with older devices; and security and privacy issues. It is unclear at this point how suitable 5G technologies will be for IoT / M2M deployments in the short term, and 4G variants are expected to maintain the IoT deployment stronghold for the next several years.



CELLULAR FALLBACK

During the early phases of any new cellular generation deployment, it often is the case that the newer generation is not fully deployed everywhere.

Typically, the geographical coverage starts small and expands over time. Thus, the cellular devices must support multiple generations of technologies until coverage is fully complete for the new technology.

Cellular radios essentially “fallback” from newer generations to older generations when the newer generation service is not available at a particular geographical location. The control of when to fall back (including which technology to fall back to) is incorporated in the Subscriber Identity Module (SIM) or other radio firmware.



Two Fallback Mechanisms

To accommodate this fallback requirement in GSM, all 3G cellular devices—modules, smartphones, and cellphones—are expected to function in 2G GSM / GPRS and EDGE modes. This allows them to be used in areas where 3G UMTS service may not be available. This increases the cost of the cellular device, but is an acceptable trade-off since it is essential to provide robust service coverage for all users of the services.

Similarly, in CDMA, the 3G EV-DO modules, smartphones, and cellphones are capable of being used in 2G 1xRTT modes, thereby enabling use in markets where 3G may not be available (this, however, is a relatively rare situation).

In 4G LTE, there are two technology fallback mechanisms. For the CDMA operators who are deploying LTE, the radio must fall back from LTE to EV-DO and 1xRTT. For the GSM operators deploying LTE, the radio must fall back from LTE to UMTS (HSPA) and then to EDGE or GPRS (since 3G is not available everywhere).

LTE-Only

These fallback mechanisms increase the complexity and cost of the chipsets within the current modules and smartphones. In time, when LTE is commonly available everywhere that cellular services are deployed, it makes sense to use radios that only use LTE services—called LTE-Only modules. These have begun appearing for purchase, and more manufacturers will deploy LTE-Only modules soon.

LTE-Only can reduce the cost of modules substantially. With scale, these LTE-Only devices will become less expensive than the lowest-cost 2G GPRS radios available today. In a few more years, this should be true for all suppliers that provide IoT modules.

Customers who want to migrate from 2G to 3G services to 4G may find it worthwhile to use LTE-Only modules to make the transition. This transition date is dependent on the customer's product longevity requirements.

HOW TO DETERMINE LOCATION

For many IoT applications, knowledge of the physical location of the devices is important—not only to the device but also to the application servers that process data from the devices.

For example, in consumer automotive IoT applications, knowledge of the exact location, to a reasonable degree of accuracy, of a vehicle crash is vital so that emergency first responders can be sent to the crash site quickly. Seconds may matter.

In truck telematics, a dispatch service may need to know the location of the vehicles in its fleet to optimize the selection of the correct vehicle to handle the specific event—perhaps it is the nearest vehicle to the pickup or has the available cargo capacity for the job. In both cases, knowledge of device location is important to a particular degree of accuracy.

The E911 location accuracy requirements are not necessarily sufficient for some IoT applications.

For emergency dispatch, the U.S. Federal Communications Commission (FCC) has defined location accuracy requirements that must be made available to Public Safety Access Point (PSAP) personnel. These often are called the “E911” requirements, since the number 911 is used to access emergency services from landline phones and cellphones.

The E911 accuracy requirements are not necessarily sufficient for some IoT applications. The location error may not allow proper calculation of routes or dispatch with sufficient optimization. **For these applications, more accurate location fix mechanisms must be used.**

Location from Cellular Network

Location-Based Services—To support the E911 requirements for physically mobile cellphones used by humans (i.e., which are not fixed at a particular address like a landline phone), cellular operators have implemented various device location mechanisms in their networks. These generally rely on classic radio triangulation techniques that provide the specified degree of accuracy for the E911 requirements.

These network-based location fixes are made available to the PSAP personnel as needed, and also are available from operators as Location-Based Service (LBS) information, generally for a fee charged for each location fix of a cellular device. Unfortunately, the cost of these location fixes may be too high for many IoT uses, and the accuracy may not be sufficient for some uses and, so, has not proven to be a common technique.

Therefore, using the U.S. GPS (and the other Global Navigation Satellite Systems or GNSS) system may well prove to be a superior solution for most IoT applications.



GLOBAL POSITIONING SYSTEM

Many cellphones now are equipped with Global Positioning System (GPS) support that allows the phones to determine their location and provide that data to the cellular network, for E911 and other purposes. Enabling this function often is an available choice in cellphones equipped with GPS.

In IoT applications, most modules have built-in GPS support (sometimes including support for both systems operated by the U.S. and Russian governments). In the future, support for the European Galileo and other national satellite systems (collectively called Global Navigation Satellite Systems or GNSS) will be implemented in many modules and handsets.

These can be used by the application firmware in the device, when needed, for a particular function, such as responding to a location fix request by a dispatch application.

In the latter half of the last century, the U.S. Department of Defense deployed a set of 24 satellites into Earth orbit for a very singular purpose—it allowed a GPS-equipped device to determine its location anywhere on the planet with very good accuracy.

Originally intended for military uses, the U.S. government made the system and its information available for civilian use in the 1980s, without any fee or subscription charge. This enabled a large number of new location applications around the world.

For example, the truck telematics industry relies heavily on GPS to locate trucks and trailers. Hikers and off-road personnel use hand-held GPS trackers to avoid becoming lost. High-accuracy augmented GPS services are used by farmers to locate the farm machinery to within a few centimeters of present location. Survey equipment can use GPS to accurately measure locations for mapping and thereby increase map quality and improve route guidance systems in vehicles.

A satellite service similar to U.S. GPS, called GLONASS, has been deployed by the Russian government. The European Union is in the process of launching its own system called Galileo. The complete 30-satellite Galileo system (24 operational and six active spares) is expected by 2020.



The Indian government has introduced its own localized system, called IRNSS, to determine location, but only over the Indian subcontinent. Similarly, the Chinese government satellite location systems, called BeiDou-1, BeiDou-2, and BeiDou-3, for global location coverage similar to the U.S. and Russian systems, with the final launch occurring in early 2018.

In time, Galileo will provide a free, low-precision location fix with an accuracy of one meter, with higher precision fixes provided for a fee. Since it is a new system, it also has new features that are not available in the older U.S. GPS and Russian GLONASS systems. For example, Galileo has radios that will support a unique relay service for Search-and-Rescue (SAR) distress signals, allowing emergency dispatch around the planet.

In addition to the GNSS transmissions, enhancements are available to dramatically improve location accuracy. For example, a set of ground-based references can be used by certain receivers to greatly enhance the basic accuracy of the U.S. GPS system from 15 meters to less than 10 centimeters. This enhanced system is called Differential GPS and enables users to employ automated equipment that need a very high accuracy location fix.

How Does Basic GPS Work?

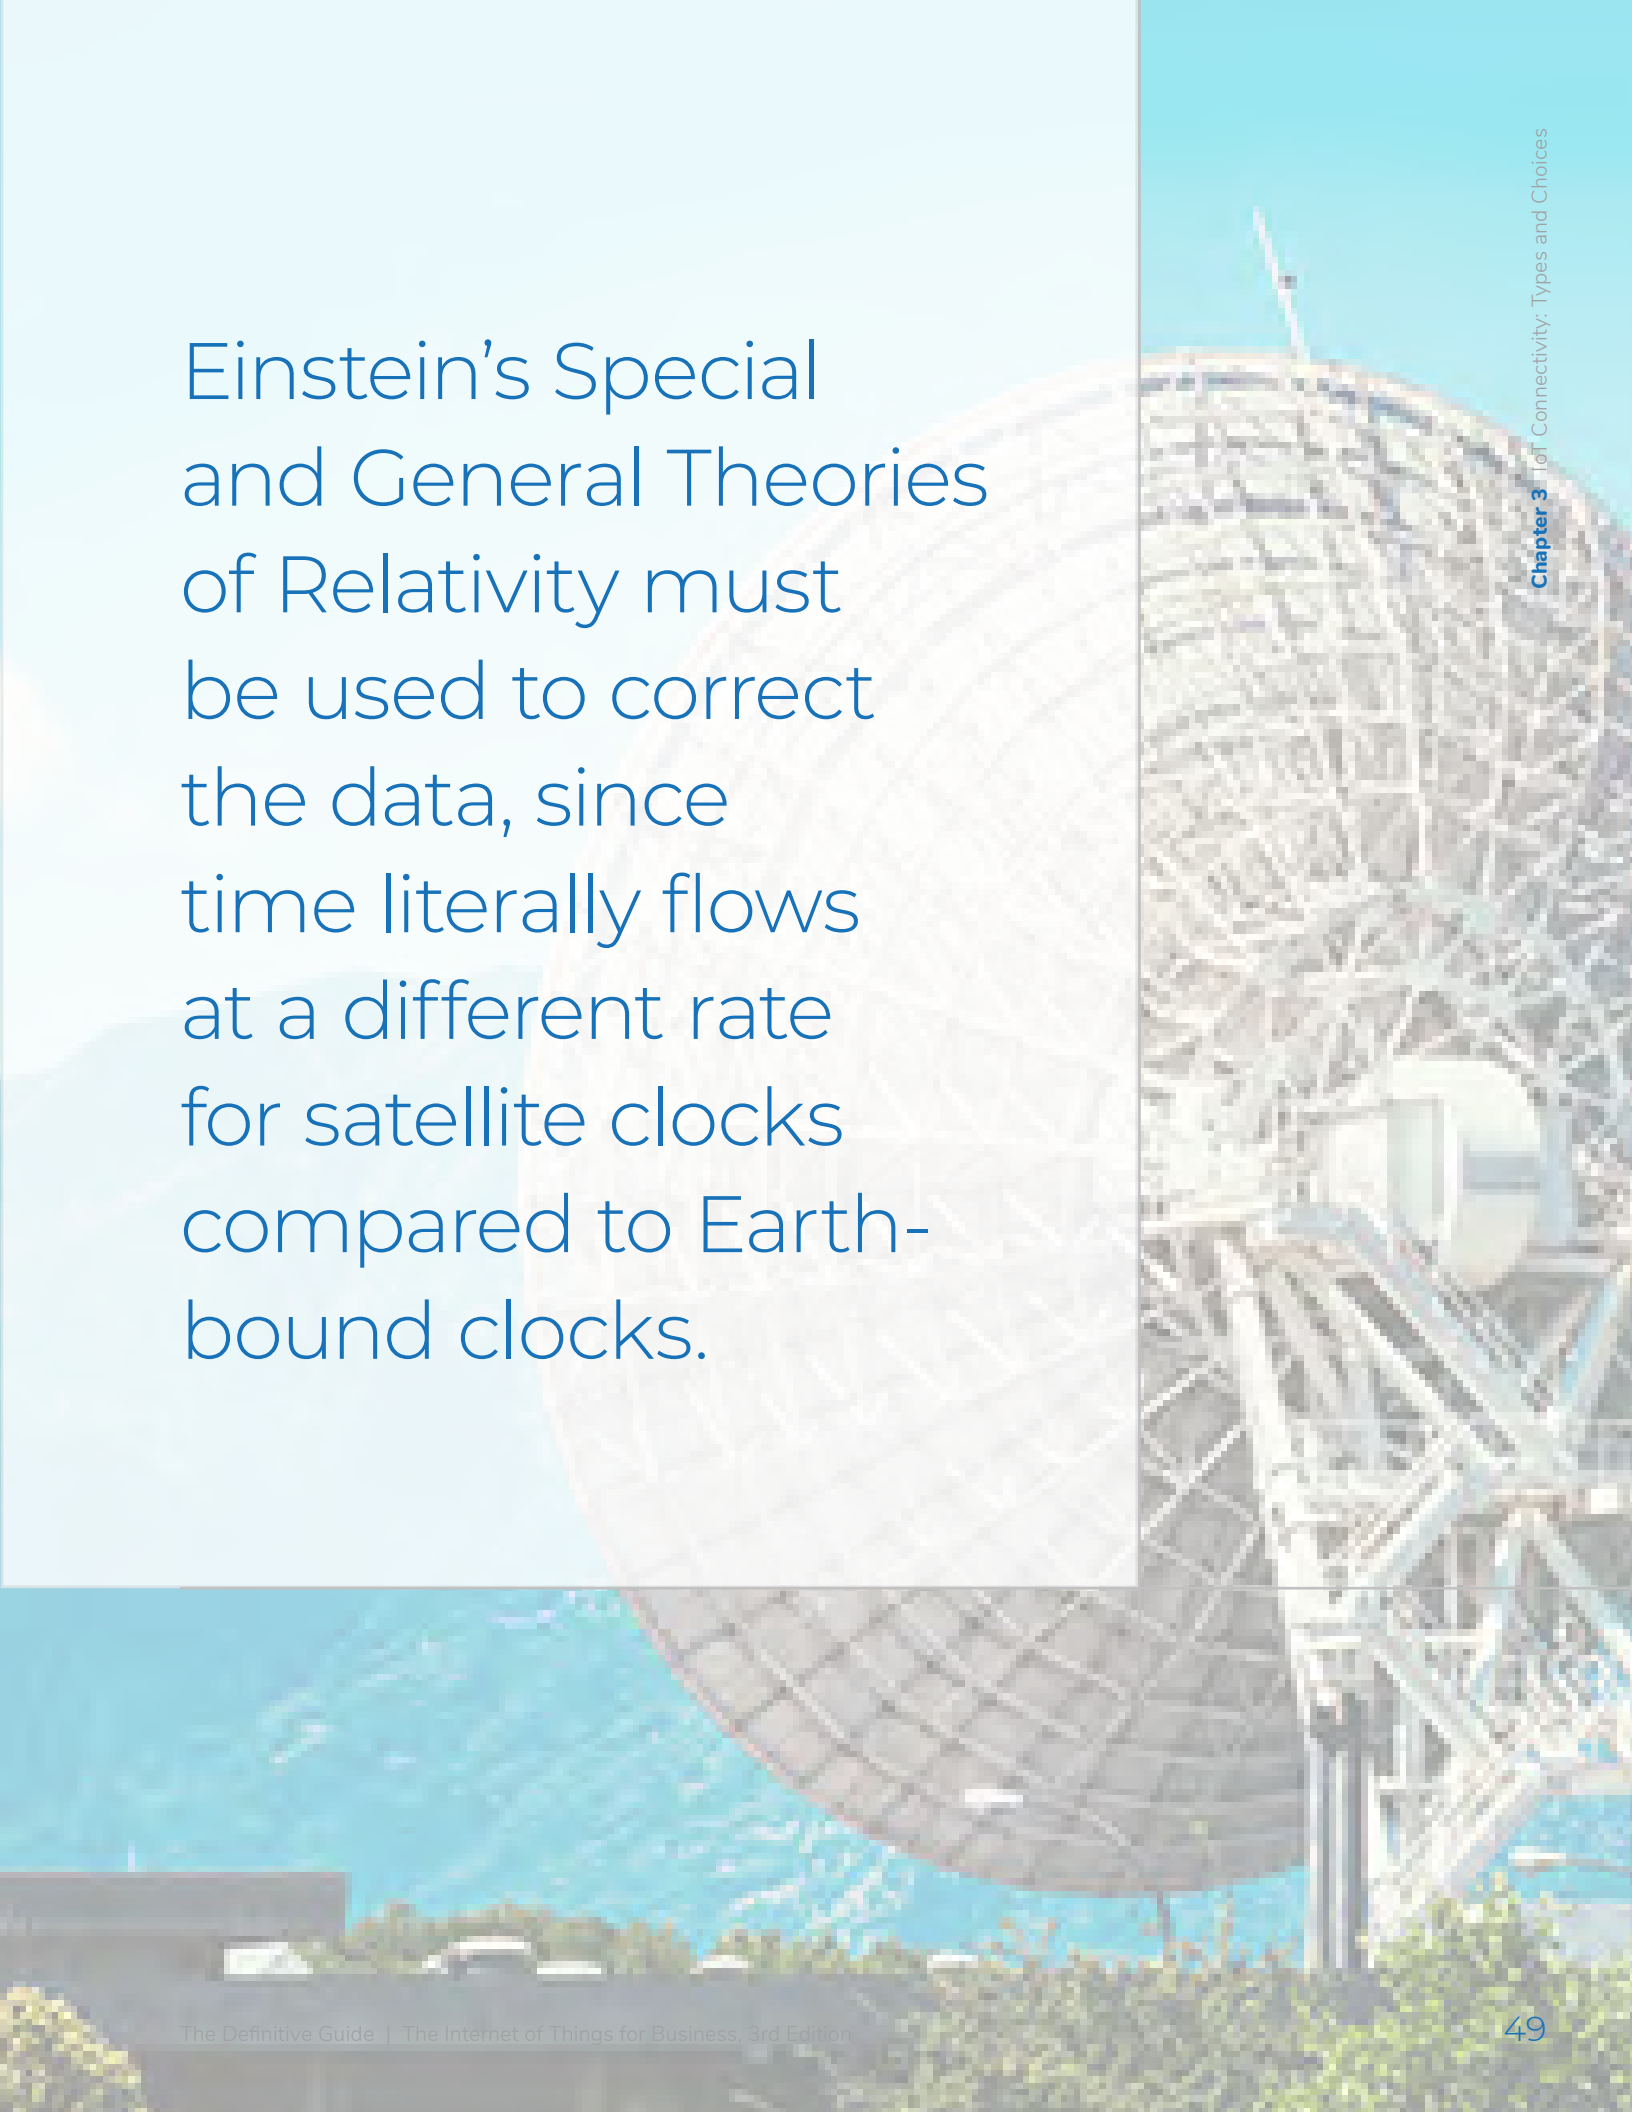
In the U.S. GPS system, more than 30 GPS satellites orbit the Earth twice a day, in a very precise manner, at an altitude of approximately 20,000 km while transmitting accurate time signals from its on-board atomic clocks to ground GPS receivers.

These GPS receivers take the received time data and use triangulation (more correctly, “trilateration” using points of intersection of circles on a sphere; angles are not measured) techniques to determine the location of the receiver. The receiver essentially compares the time a signal was transmitted by a GPS satellite to the time it was received. This time difference allows the receiver to determine its distance from that satellite.

When time difference and distance are determined from a number of GPS satellites, the location of the receiver can be determined within 5 to 10 meters of accuracy on the surface of the Earth. At least three satellites must be used for a latitude-longitude fix on the surface of the Earth, and a fourth satellite then can determine the altitude of the receiver.

It should be emphasized that the above is a very general description of the method used to determine location from the GPS satellite signals. There are a number of other factors that affect the accuracy and are taken into account by sophisticated receivers.

Ground-based references can be used by certain receivers to greatly enhance the basic accuracy of the GPS system from 15 meters to less than 10 centimeters.



Einstein's Special and General Theories of Relativity must be used to correct the data, since time literally flows at a different rate for satellite clocks compared to Earth-bound clocks.

For example, the more satellites the receiver listens to, the better the accuracy. Thus, a 10- or 12-channel GPS receiver (which allows it to listen to 10 or 12 GPS satellites simultaneously) generally will provide a more accurate location fix than an older 4- or 6-channel receiver. Modern GNSS chips can listen to more than one satellite system simultaneously (such as GPS and GLONASS) for best accuracy.

Furthermore, since the GNSS satellites are in motion and are quite far above the Earth's surface (i.e., operating in reduced gravity), Einstein's Special and General Theories of Relativity must be used to correct the data, since time literally flows at a different rate for satellite clocks compared to Earth-bound clocks.

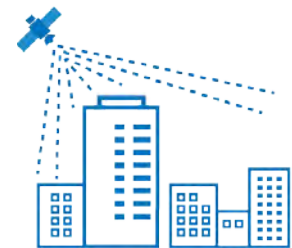
Without proper compensation, the relativistic effects of the speed of the satellites, combined with their height, could create a net error of about 38 microseconds per day for the satellite clock compared to an identical ground-based clock. This may seem quite inconsequential, but the difference in time can make the location fix inaccurate within a matter of minutes, to beyond the 5- to 10-meter accuracy of the system. Then, accumulated errors could make the location fixes completely unreliable and unusable in a matter of days since GPS requires nanosecond time accuracy.

Fortunately, the GNSS receivers use these Einsteinian Relativity calculations and corrects to ensure that the time and location accuracy is excellent, and remains excellent, under most conditions.

With multiple location fixes (i.e., taken over time), the data fixes also can be used to determine other information, such as speed and direction (i.e., velocity). Sophisticated tracking devices can use the data to display the location and provide route guidance in ways friendlier than a simple latitude-longitude-height-time record that is displayed on a graphical moving map, for example.

Limitations of GNSS


Location fixes from global location devices are not perfect. In "urban canyons" (i.e., within cities with tall buildings), it may be difficult for GNSS receivers to lock onto more than a few satellites since the signals may be blocked by buildings. This may reduce the accuracy substantially. Regardless, it may remain sufficiently capable for many uses of that location data. A higher-performance GNSS receiver with many channels may perform better in urban canyons since it has a better chance of listening to satellites that may be "visible" and not blocked by tall buildings.



The signal strength is low enough that many GNSS receivers cannot listen to the signals from the satellites when inside buildings and underground garages. This limits use in indoor applications, although new technologies for indoor tracking to supplement outdoor GNSS performance are in the works.

Sometimes, heavily overcast days can reduce the strength of the GNSS signal, enough to prevent the receiver from locking on to the signals, particularly when the receiver has been re-started from a power-off condition. If the internal clock of the receiver is not sufficiently accurate, the measured time may have drifted and the receiver could be attempting to listen to a set of satellite signals that are not present. Those particular satellites may not be visible. Under these conditions, it may take a while for the receiver to lock onto the satellite transmissions and provide sufficient accuracy.

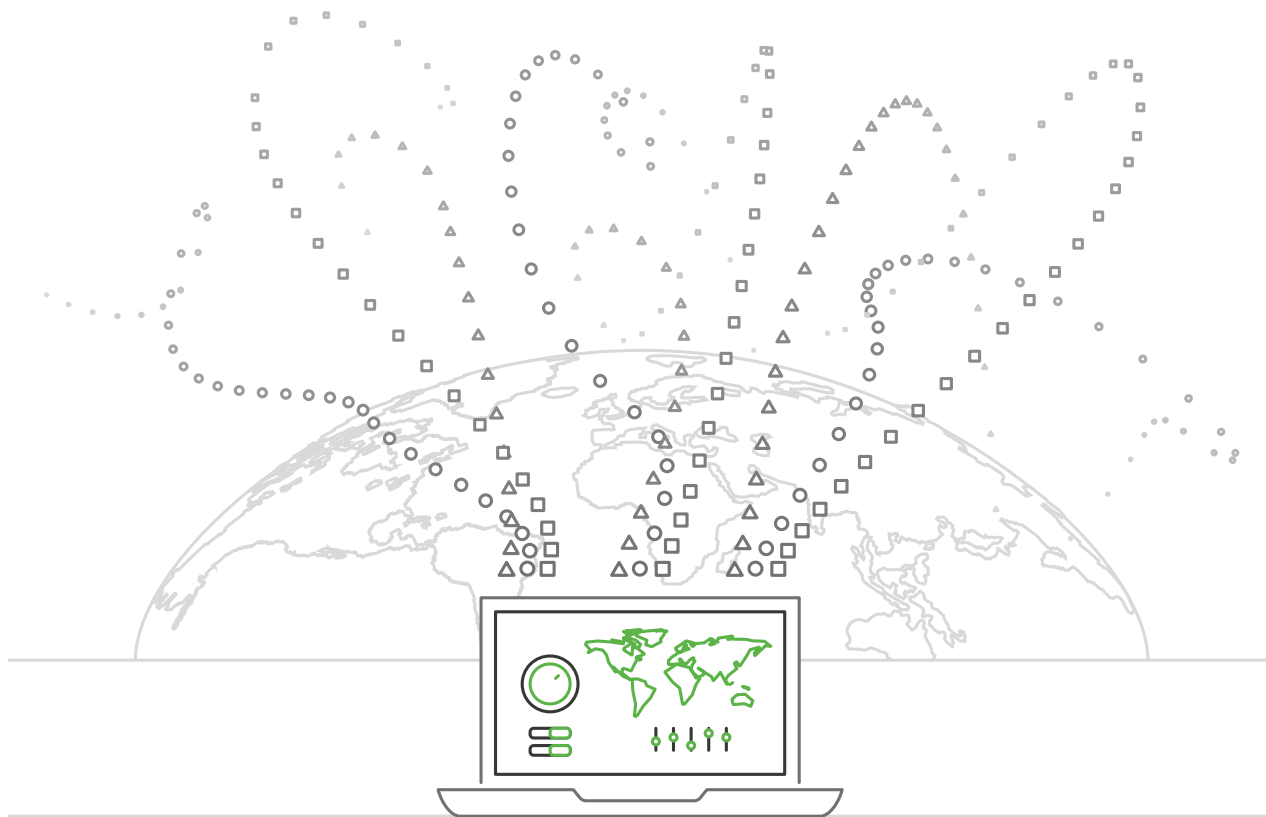
IoT applications that depend on location accuracy must take this potential for incorrect location fix data into account when deploying a device in the field.



CONNECTIVITY MANAGEMENT PLATFORMS

- 54 WHAT IS A CONNECTIVITY MANAGEMENT PLATFORM?
- 55 THE DIFFICULTIES OF MANAGING IoT CONNECTIVITY
- 56 WHY BUSINESSES NEED CONNECTIVITY
MANAGEMENT PLATFORMS
- 59 ESSENTIAL CONNECTIVITY MANAGEMENT
PLATFORM FEATURES
- 61 CMPs ARE INTEGRAL TO
THE IoT ENVIRONMENT





CONNECTIVITY MANAGEMENT PLATFORMS

Advanced, diversified, and cost-effective connectivity is integral to the success of Internet of Things and machine-to-machine communications. Smart devices and network endpoints generate unprecedented amounts of data that must be collected, stored, and analyzed to perform IoT-driven business operations and services. These processes involve data transmission using various connectivity services and technologies that organizations must manage effectively to maximize the value potential of their IoT deployments.

Furthermore, as wireless technologies become commoditized and new technologies are being introduced regularly, successful IoT deployments are able to rapidly change with the evolving wireless landscape to take advantage of new technologies and reduced rates. The result is an environment where multiple technologies are utilized across multiple suppliers, adding operational complexities and costs to IoT deployments.

Connectivity Management Platforms (CMPs) exist to help reduce the complexities of managing an IoT deployment, but also can make the problem worse if not utilized correctly. This chapter explores the role of CMPs in the modern IoT infrastructure, IoT organizations, and the evolving IoT enterprise landscape.

WHAT IS A CONNECTIVITY MANAGEMENT PLATFORM?

IoT devices have evolved into smart network endpoints that extend the reach of cloud operating systems and perform intelligent actions on their own. These devices don't require unique software embedded in every hardware device.

Large IoT networks leverage multiple connectivity providers to address diverse business needs. The result is an increasingly complex network of smart devices that must be managed as a ubiquitous system. These complexities not only cause problems from a technology perspective but may drain financial and management resources merely to keep IoT systems operational as a unified network.

With limited resource availability and increasing complexity of IoT networks, organizations must automate the way they manage, configure, control, and track IoT device information. CMPs automate these processes to enable effective deployment, management, and utilization of IoT networks that span disparate geographical locations, connected with multiple service providers, and designed to scale exponentially.

THE DIFFICULTIES OF MANAGING IoT CONNECTIVITY

Managing IoT deployments on any wireless technology is a complicated endeavor. Consider the geographic considerations of IoT networks that span multiple countries, each presenting its own set of financial, legal, compliance, and technology challenges.

Lack of visibility and control is inherent in these circumstances, especially when operational excellence of multinational organizations is tied with supply chains, inventories, logistics, and departments located in globally dispersed sites—all of which use connected systems and devices to operate.

Merely keeping IoT deployments connected is a challenge in itself. Even with high signal strengths, IoT networks may be impacted with hardware, firmware, configuration, or application-level issues. Fast and effective issue resolution is dependent upon real-time monitoring and accurate issue-tracking capabilities. Failure to resolve connectivity issues not only increases operational costs but also risks system-wide outages and downtime that may lead to non-compliance and legal implications, as well as customer dissatisfaction and damaged brand reputation.



In today's wireless landscape, organizations need multiple connectivity options to support large-scale IoT deployments. No single technology or provider delivers the most reliable, cost-effective, high quality, and advanced connectivity services. This means organizations must manage multiple supplier agreements and manage their deployment across multiple management platforms. And each platform will have its own set of features, functionality, and capabilities. Disparate platforms force organizations to collect, standardize, and analyze data from various platforms to perform desired IoT operations. This inability to standardize connectivity operations prevents cost optimization and restricts scalability.

Organizations pursuing scalable IoT solutions need to find ways to create a single interface to proactively manage and monitor their IoT deployments, connectivity channels, and the other associated financial, legal, and technical aspects. This capability ensures reliable and high-value performance of IoT systems—regardless of deployment location, network span, or operator diversification.

WHY BUSINESSES NEED CONNECTIVITY MANAGEMENT PLATFORMS

The Internet of Things presents vast strategic business advantages for the modern enterprise. IoT brings automation and intelligence to everyday objects, devices, and things—looking to revolutionize the consumer and enterprise market segment alike.

And mirroring this automation and intelligence with a cost effective and efficient management infrastructure has an often overlooked impact on determining the success or failure of an IoT solution. Managing IoT deployments cost effectively is essential, and exponentially difficult with a multi-technology, multi-provider deployment.

A centralized CMP is required to address the consolidation challenges associated with operating large deployments of IoT devices across multiple providers and technologies. The following areas will be key for any business looking to gain value from a CMP in a highly scalable IoT environment.

Market Drivers for CMPs

The economies of scale for using IoT platforms will affect all CMP consumers, from the connectivity provider to the OEM / application service provider to the developers. Scalable, flexible CMPs will be essential for these partners to transition from custom vertical solutions to horizontal platform-based solutions.

This has the main advantage of enabling them to work with high velocity, volume, and variety—or the “three Vs”. High velocity in a CMP will spread the cost of infrastructure across many industry verticals. High volume will result in shorter design cycles and rapid deployments. And variety as part of a CMP means integrated network information can enable a richer application experience.

It also is crucial to recognize the diversity in connectivity options as a primary driver of the need for CMPs. Enterprises will first need to consider whether they want wireless or wired connectivity or, for some applications, they may need both. Within the realm of wireless, options range from varieties of cellular and satellite data services to short-range wireless technologies such as Wi-Fi and personal area networking, including Bluetooth or ZigBee. Best-in-class CMPs should be able to deal with all of these technologies and manage the protocols effectively.

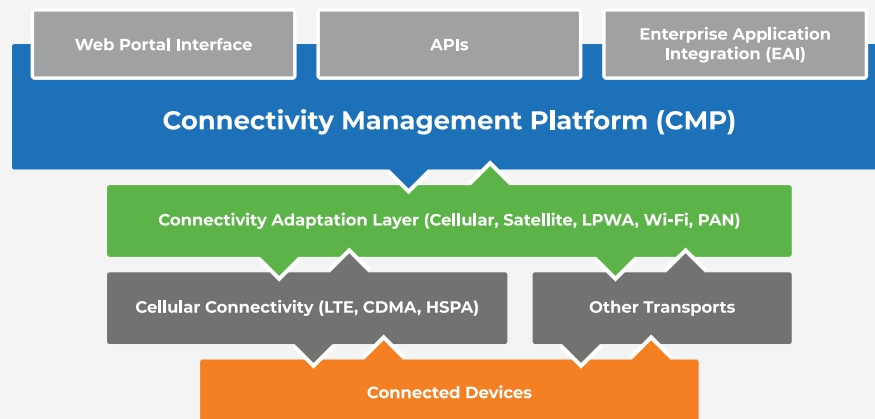
CMP Technology Needs

The most integral connectivity management functions sought by IoT connectivity operators, as well as enterprise consumers and developers, range from network visibility and control to data security to end-user management, billing, and reporting.

Diverse vertical applications and environments will need to access a typical CMP for industries such as healthcare, automotive, and manufacturing. Ease of integration with enterprise applications, including APIs, network data feeds, etc., is essential.

CMP Architecture

The CMP plays a vital role in an IoT technology stack, enabling the promised technology and business functions within the IoT infrastructure. Enterprise users expect “single pane of glass” access for multiple connectivity technologies, whether they’re using cellular or Wi-Fi, so the CMP must support a diverse set of IoT applications.



CMPs can offer a diverse range of capabilities to serve your organization's unique requirements. Getting the right fit is critical for effective connectivity management.

ESSENTIAL CONNECTIVITY MANAGEMENT PLATFORM FEATURES

CMPs can offer a diverse range of capabilities to serve your organization's unique requirements. Getting the right fit is critical for effective connectivity management.

While not every solution in the market addresses all of the fundamental enterprise and operator needs, you'll want to consider these factors in your prospective CMP:



Core CMP Functions—Your CMP should enable a range of business and technology functions associated with the connectivity of your IoT endpoints. These include the management of zones, devices, users, accounts, pricing, policies, billing and invoicing, alerts and reporting, and SIM lifecycle management, among others.



Global and Multi-Carrier Footprint—IoT networks that span disparate geographical locations and multiple countries use services from multiple providers. Your CMP must support this capability.



Security—Your CMP should offer features that ensure high data security, availability, and end-user privacy. These features typically include identity and access management, audit trails, anomaly detection, denial of service prevention, and strong encryption.



Multiple Carrier / Multiple Technology Support—Your CMP must support multiple connectivity technologies—from cellular (LTE, GSM, HSPA, CDMA) to non-cellular (Wi-Fi, Bluetooth ZigBee, satellite, and LPWA). And you should have the flexibility to bring your own connectivity from the operator or carrier of your choice. All types of data transmission protocols, technologies, and standards should be supported across different transmission layers as required by your IoT solutions.



Quick Time to Market—It takes long design and deployment cycles before custom vertical solutions are released and create value for your IoT service organization. Using a CMP ensures that IoT connectivity is managed right from the beginning, thereby making the entire getting-to-market process much simpler and quicker.

It is crucial to recognize the diversity in connectivity options as a primary driver of the need for CMPs.



Flexible Pricing—Enterprises need to accommodate a range of pricing from multiple connectivity providers. To reduce the total cost of connectivity services, your CMP should help you optimize pricing for different IoT app needs in different business functions and spread the infrastructure cost across multiple tenants and verticals.



Billing and Real-Time Data—Accurate billing information updated in real time ensures that you only pay for what your devices consume.



Variety of IoT Application Needs—Different IoT applications present different needs for connectivity management. Your CMP should support these diverse and evolving needs as more functionality is added to your IoT applications.



Open System—Strong integration with your infrastructure and other enterprise applications via open APIs.



Flexible Accounts and User Access—Progressive and agile organizations scale teams rapidly to meet evolving resource requirements. The ability to create and manage new accounts and streamline user access enables effective connectivity management.




Diagnostics and Fault Resolution—CMPs can extend the monitoring and diagnostic capabilities of your IoT systems with rich graphical representations of connectivity patterns and defined performance metrics.



Light Touch Integration—Your CMP should operate as a standalone tool and connect with other enterprise applications only as required.



Highly Scalable and Elastic—Your IoT needs will change. Consider a connectivity management solution that can accommodate changes. You want a solution that is future proven so as to withstand changes in the connectivity marketplace.



Without a single-pane-of-glass view for managing multiple connectivity services across different IoT deployments, organizations risk connectivity issues that present severe business, technology, and legal implications. The traditional practice of using manual processes or individual platforms to manage IoT connectivity service is both time consuming and ineffective, but CMPs can help solve this issue.

CMPS ARE INTEGRAL TO THE IoT ENVIRONMENT

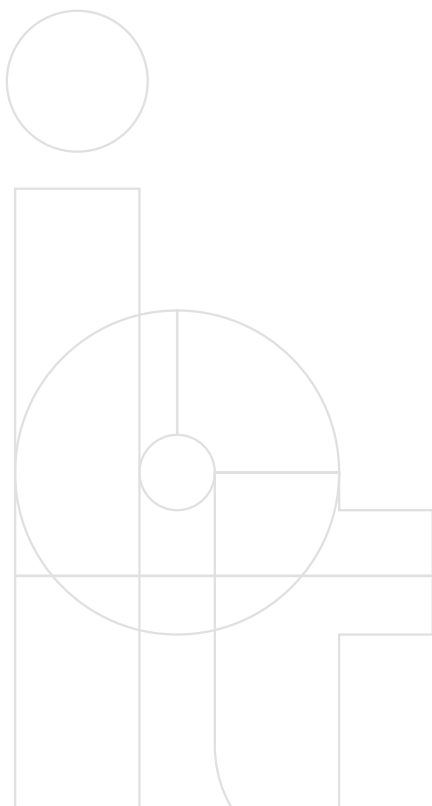
IoT deployments are driving business model innovation among progressive organizations pursuing smart technologies and advanced connectivity solutions to generate unprecedented new revenue streams.

IoT deployments are driving business model innovation among organizations pursuing smart technologies and advanced connectivity solutions to generate unprecedented new revenue streams. Intelligent IoT devices generate invaluable data in real time that is transmitted to back-end systems through various transport systems, including cellular and Wi-Fi communication. Today's agile organizations increasingly rotate data between different processes, functions, and teams.



Advanced IoT connectivity systems that must scale rapidly to meet varying organizational demands tend to create a performance bottleneck among agile enterprises if users, technologies, performance, and processes are not managed effectively. These issues only will grow with the rapid adoption of IoT technologies. And since every company needs to be technology aware in the present IoT-driven enterprise landscape, organizations will pursue a diverse range of connectivity solutions to serve their specific exploding IoT demands.

The bottom line—effective management of IoT connectivity can be a major contributing factor to the success of an IoT deployment. Managing multiple technologies and multiple providers from a single interface enables scalability, reduces operational complexities and costs, and enables future product evolution with minimal impact to business processes.



IoT SENSORS AND DATA COLLECTION

65

WHAT IS A SENSOR?

66

SENSOR TYPES

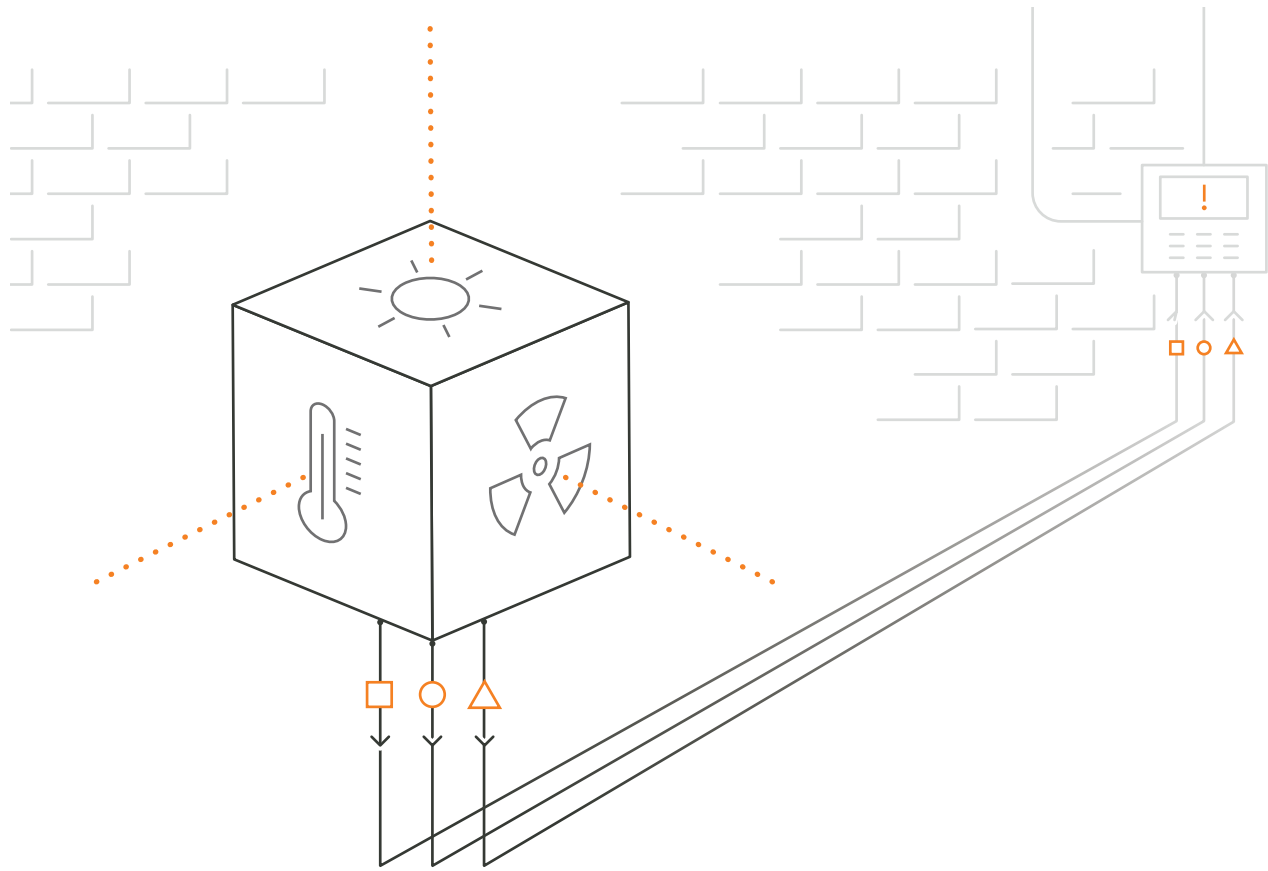
72

CONVERSION TO DIGITAL DATA

76

CALIBRATION AND LINEARIZATION





IoT SENSORS AND DATA COLLECTION

When deploying Internet of Things application devices, the connected device generally needs to report more than just its physical location (although that is a very typical use).

In this chapter, we describe a few of the more common sensor types, what they measure, and how to use them.

WHAT IS A SENSOR?

Dictionary.com defines a sensor as “a mechanical or electrical device sensitive to light, temperature, radiation level, or the like, that transmits a signal to a measuring or control instrument.”

This is an easily understood definition that can be reworded a bit in the context of IoT applications: a sensor is “a device, generally small and mechanical, that is sensitive to a measurable physical parameter and provides a measurable signal level directly related to the measured amount of that physical parameter.”

For example, an IoT device may measure a particular physical parameter, such as temperature, at a location for an application purpose. These physical parameter measurements require sensors that are capable of measuring, recording, and transmitting the specific value of that physical parameter for the IoT application to fulfill its functions.

Sensors often are integrated circuits that are designed for these kinds of IoT applications since their small size and low cost make them appropriate choices. For example, many of the sensors described in this chapter are available in high-end smartphones. These include accelerometers, thermometers, gyroscopes, magnetometers, and heart rate monitors, just to name a few. But there are other sensors that are unique to a particular industry or market.



SENSOR TYPES

In most typical sensors, the specific mechanism used to measure the physical parameter depends on the parameter ranges being measured, the desired sensitivity and accuracy, whether the sensor is exposed to adverse environmental conditions, the cost target, etc.

Since it is nearly impossible to list every possible sensor, its type and purpose, its capabilities, and the physical parameter that is measured, this section focuses on general descriptions of a few types of sensors rather than making specific recommendations.



Accelerometers

Acceleration is a change in velocity (a change of speed and / or direction). Accelerometers are devices that measure acceleration. The parameter being measured may be a static force, such as gravity exerted on a device. Other accelerometers make dynamic force measurements so as to measure motion changes and vibration.

An example of an accelerometer is a chip in a moving vehicle that measures changes of speed and uses high acceleration (deceleration) readings—such as those experienced in an automobile accident—to trigger airbags to protect the passengers.

In some industrial applications, vibrations detected by an acceleration sensor could be an excellent indicator of a potential problem with a moving part, such as a motor with bearings that are worn. Timely transmission of data from vibration sensors can enable early detection of potential problems where preventative maintenance could avoid catastrophic failures.



Multi-Axis Accelerometers and Sensitivity

In some applications, there is a need to measure the change in speed, or vibration, in more than one direction (or dimension). Thus, some accelerometers can take readings from more than one axis. Typically, a two-axis sensor measures motion changes and vibration in two dimensions, and the third axis on a three-axis sensor can provide information for three-dimensional physical motion detection.

Accelerometers use differing techniques for measuring the actual motion changes. Generally, there is a physical component that changes an electrically measured characteristic (such as capacitance or resistance) in a material when changes in motion are detected.

Due to the types of accelerations being measured, the sensitivity of the accelerometer often is in a limited range to the required accuracy specific to a particular application use. The choice of sensor, therefore, depends on the specific range of acceleration values to be measured for that application.

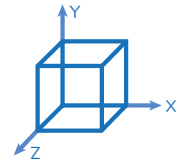
For example, a shock sensor designed to release an airbag in a vehicle accident measures quite a different range of acceleration compared to a vibration sensor that measures vibration on a high-speed motor to monitor its bearings. The sensitivity, range, and accuracy required for these widely disparate applications is, naturally, quite different.

Temperature Sensors

Temperature is a physical parameter that often is measured and reported—particularly in industrial applications where an accurate temperature reading may be needed for process control. Depending on the desired range, there are various types of available sensors for measuring temperature.

Silicon chip (semiconductor) sensors are used in the range from -50 to +150 degrees C (Celsius scale). These are quite accurate and linear—to within 1 degree C—without a need for extensive calibration. They are as rugged as most integrated circuits (plastic package and metal can style) and relatively inexpensive.

Thermistor sensors can cover a wider range—from -100 to +450 degrees C—for covering more applications. A thermistor often is more accurate than a silicon chip temperature sensor, albeit at a slightly higher cost per sensor. More important, they require complex correction algorithms to achieve that good accuracy and linearity over the desired temperature range.



Resistance Temperature Detectors (RTD) provide yet more range, from -250 to +900 degrees C, but are quite difficult to use since they are more fragile than other types of temperature sensors. They are the most accurate—often a hundred times more accurate than a silicon chip sensor—although this requires the same complex solutions for linearization as thermistors, and some models can be quite expensive.

Finally, for the widest temperature range, particularly for very high temperatures, a thermocouple is the correct choice. They are quite rugged and can be used from -250 to +2000 degrees C for many industrial applications, such as chemical process monitors and high-temperature electric furnaces (for example, those used in the semiconductor industry).

One important fact about temperature sensors is that the response time for measuring changes in the temperature data can be quite slow since temperature changes are not as “rapid” as other measured physical parameters. The sensor readings must settle and equalize to the temperature being measured. This must be taken into account when taking readings.

Light Sensors

Light sensors cover a broad range of potential applications, from automated brightness control in cellphones to medical diagnostic equipment. Not surprisingly, there also is a wide range of available light sensors that use different methods for measuring light.

A very early example of ambient light sensors used in consumer applications are photocells within lamps that automatically turn the lamps on at dusk and turn them off at sunrise. These are simple ambient light detectors, with equally simple sensitivity controls that are adjusted manually by the owner of the product. The actual value (in lumens) of the ambient light is not measured or reported—it simply is used to perform the designed function of automatically turning the lamp on and off.

Simple light sensors also can be used for proximity detection. Counters in manufacturing systems use the presence or absence of light on photocells to measure products being moved past the counter on conveyors. Closing garage doors can reverse direction to avoid hurting children or pets that cross under the door by sending a beam of light across the door opening to a photocell. Once the beam is cut, the sensors send a message to reverse the garage door direction.

Light sensors
can be used
with light
that is not
visible to
human eyes.



Often, light sensors can be used with light that is not visible to human eyes. Infrared light sensors can be used as motion sensors in alarm systems or to automatically light a driveway or passage when people and pets come into range. Full-range light sensors are used when the light measurements need to correspond to human vision.

As with other types of sensors, the mechanism used to measure ambient light varies depending on the application. Simple Cadmium Sulfide (CdS) or Cadmium Selenide (CdSe) photoresistors change their resistance as a function of the ambient light. This resistance change can be measured by electronic circuits to provide an indication of a change in the ambient light. It should be noted that these photocell devices can be significantly affected by temperature and are quite unsuitable when accuracy is required.

Common uses of photoresistors include automated light controls in lamps, dimmers in alarm clocks and audio system displays, or control of street lighting systems—where the accuracy of the reading is not of paramount importance.

Photodiodes and phototransistors with active semiconductor junctions are used when greater accuracy is required, since the ambient light is converted into a measurable current that can be amplified or converted for a measurement. This measured current can be used to determine the amount of ambient light on the sensor. And since semiconductor junctions are affected by ambient light, integrated circuits where this effect is not desired must be enclosed in opaque packages.

MEMS Sensors

In modern, high-end smartphones, integrated chip sensors that measure motion, direction, pressure, magnetic fields, rotation speeds, and more are becoming quite common. These can be used to augment the location information and human motion in the cellphone.

In chip form, these usually are Micro-Electro-Mechanical Systems (MEMS) sensors for many different parameter measurements. The implementation of MEMS uses ultra-miniaturized physical structures—beams, arms, and associated electronics—to measure the motion of the structures when the chips move. The physical motion is converted into electrical signals that can be measured for the specific function being measured—for example, whether it is rotational motion or air pressure. The sensor essentially converts a mechanical motion into an electrical signal.

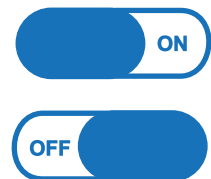
A gyro sensor, for example, senses rotational motion and changes in orientation. These can be used in a variety of applications, such as correcting for hand-held shake in video and still-image cameras or for motion sensing for video games. In smartphone applications, a screen display can be rotated automatically between portrait to landscape display modes when the phone is physically rotated.

MEMS sensors generally are manufactured in the same large-scale facilities as semiconductors or chips. This means that the mechanical precision of the devices can be very high and allow for excellent, reliable performance at low cost.

Simple Switch Sensors

At the very low end of the sensor markets are the simple state or position sensors that provide an “open” or “closed” state information. A door or window sensor used in security systems often is a simple magnetic reed-relay switch that opens, or closes, an electrical circuit depending on the position of a small magnet physically located close to the switch.

These simple magnetic reed-relay switches can be used for sensing when a cabinet—such as a medicine cabinet, oven door, or food storage compartment—has been opened in a senior citizen’s home-monitoring IoT application. A detection of the change of state of such a switch—from open to closed or vice-versa—can be interpreted as evidence that the monitored person has performed their expected regular daily routine.



Specialized Sensors

In industrial and simple applications, many standard sensors have been developed and commercialized over the years. These have evolved and improved over time. The cost and size of sensors have been reduced with increasing efficiency and practicality.

Recently, however, particularly with the start of the IoT revolution, there has been a great demand for a variety of new parameters to be measured at ever-larger scale and ever-lower cost.

The healthcare industry is among those at the forefront of this revolution. New methods to measure human medical parameters are being researched and commercialized, and this has seen an explosion of new techniques (and sensors) to measure these parameters. In medical monitoring applications, the need for new measurements, reduction in the size of devices, and the rapid adoption of wearable fitness and medical products is driving significant research and growth.

Beyond the sensors incorporated into hand-held or wearable products (e.g., smartwatches, clothing, and bracelets) and for reading basic body functions or medical monitoring products (such as continuous blood sugar monitors and insulin dispensers), there also is a need for semi-permanent sensors implanted within the human body. The research into tiny, implantable sensors has been energized by the availability of semiconductor and MEMS solutions, including for mission-critical applications such as cardiac monitoring and vision correction.

For example, medical startups are developing MEMS sensors that are implanted into pulmonary arteries using cardiac catheter procedures similar to angioplasty. These sensors can measure artery pressure and transmit the readings to a nearby wireless device within the patient's home. The readings then can be sent to a database for review by medical practitioners.

These new methods for measuring human parameters, the sensors using these methods, and devices using these sensors have been commercialized in the past few years. Newer sensors will be introduced in the next decade and will completely revolutionize the medical healthcare industry in ways that we cannot even imagine today.



CONVERSION TO DIGITAL DATA

Because of the wide variety of sensors, the types, the parameter being measured, and what physical phenomenon is converted into a measurable signal, it is difficult to provide implementation details. Thus, this section will discuss general concepts rather than specific information.

Sensors often are used in local applications, where their signal is processed using circuitry designed for that local application. However, in a sensor that is used for transmission of the measurement to remote computing and analytics systems, the measured electrical parameter must be converted into a digital value, or number, for the transmission.



Furthermore, the specific electrical signal from different sensors may vary over a wide range of current or voltage, or other electrical parameter (such as resistance or capacitance) and often must be converted and amplified into a voltage that can be measured more easily.

If necessary, the signal must be filtered electronically to eliminate signal noise or to reduce the frequency of the measurement for the requirements of the application. For example, a temperature sensor generally changes its value relatively slowly as the sensor matches its environment. Therefore, a rapid change in reported temperature may be an inaccurate reading, which should be filtered to reduce potential errors.

Device Input / Output

With devices that measure sensors for data transmissions, two input capabilities generally are available:

- A digital input pin that reports an electrical “high” or “low” value in a single digital bit (sometimes grouped into multiple pins and multiple bits).
- An analog input pin that receives a voltage from a sensor and converts that voltage, using an Analog-to-Digital Converter (ADC), to a digital number that represents the sensor value.

These devices also may have output pins where a received value is used to:

- Set a digital output pin to either a “high” or “low” state based on an instruction to do so.
- Set an analog output pin to an analog voltage, using a Digital-to-Analog Converter (DAC), representing the received digital number.

ADC Techniques

In sensors that measure parameters over a range, a single bit is insufficient—the range of the measured sensor readings must be converted into a range of digital values for the application.

However, the specific signal from a sensor may differ widely in its current or voltage or resistance value. This signal—whether it is a current or resistance change—must be “conditioned” or converted to an analog voltage. If the signal from the sensor is a voltage, it might not be in the correct range for an ADC to convert to a digital number and, thus, may require amplification to a higher or lower voltage range.

For example, a commonly available semiconductor temperature sensor provides an electric current of 1 microamp per degree Kelvin when power is applied to it. Over a useful range of -50 degrees C (or 223 degrees Kelvin) to +150 degrees C (423 degrees Kelvin), this current is approximately 223 microamps to 423 microamps.

This current can be used in a circuit with an Operational Amplifier (Op-AMP) and other components (resistors, capacitors, and diodes) to convert the current to a voltage in the desired operating temperature range being measured. This voltage then can be measured by an ADC and processed by the device taking the temperature measurement for the application function.



Some sensors, usually more complex and expensive, may have built-in functions for converting the measured physical parameter directly to a number that is sent to the device processor or communications module for transmission.

For example, a GPS device may report continuous position and time readings on a serial port using common National Maritime Electronics Association (NMEA) formats called NMEA 0183 or NMEA2000. This data already is “conditioned” into a text format that can be used by a device processor to communicate and transmit the location data.

ADC and DAC Resolution


When converting the analog voltage signal from a conditioned sensor to a digital value or number, the ADC has a pre-defined resolution based on its design. This means that the full range of the measured analog signal varies from a zero value to a maximum numerical value, with incremental steps defined by the resolution of the ADC.

For example, an 8-bit ADC will convert the voltage level to an integer between a low value of 0 to a high value of 255 to represent the value of the analog signal in approximately equal steps. This may be quite sufficient for many IoT applications.

In other applications, it may be necessary to use a 12-bit, or perhaps even a 16-bit, resolution ADC for the conversion. A 16-bit ADC provides a digital numerical value between 0 and 65535 based on the input analog voltage. With higher resolutions—particularly with low signal levels, the signal conditioning and amplification circuits may need special design to ensure that electrical noise does not result in erroneous readings.

It is important to note that resolution is not the same as accuracy or linearity. It merely identifies the number of integral steps between the lowest and the highest value being converted and reported.

A full discussion of these concepts is beyond the scope of this book. Interested readers can refer to the data sheets and applications notes from ADC and DAC suppliers for more information.



Resolution is not the same as accuracy or linearity. It merely identifies the number of integral steps between the lowest and the highest value being converted and reported.

Modules or External Processors

Quite often, the communications modules or modems—particularly cellular products used in industrial IoT applications—have multiple input and output (I/O) pins that can provide the conversions from a measured physical parameter to a number. This can be a simple on / off state using digital input pins, or an analog voltage reading with an on-board Analog to Digital Converter on an analog input pin that is converted into a number that is transmitted on the communications network.

Some modules also have digital output pins for setting a state external to the application—for example, to activate a relay to turn on a light, power on an electrical device, disable a vehicle, or perform some similar remote IoT function.

A few modules and modems also have DACs that take a digital number received from the communications network and output an analog voltage on an analog output pin in a specific voltage range. This may be used where an analog voltage is used to perhaps control the position of a liquid flow valve or the speed of a motor in an industrial IoT application.



CALIBRATION AND LINEARIZATION

As described earlier, a simple sensor application (such as a photocell that controls a lamp to turn on or off based on ambient light) may not need an accurate reading or sensor value. However, when accuracy is important for an application, calibration of the sensor signal may be needed to ensure that the data reading is accurate to the required degree.

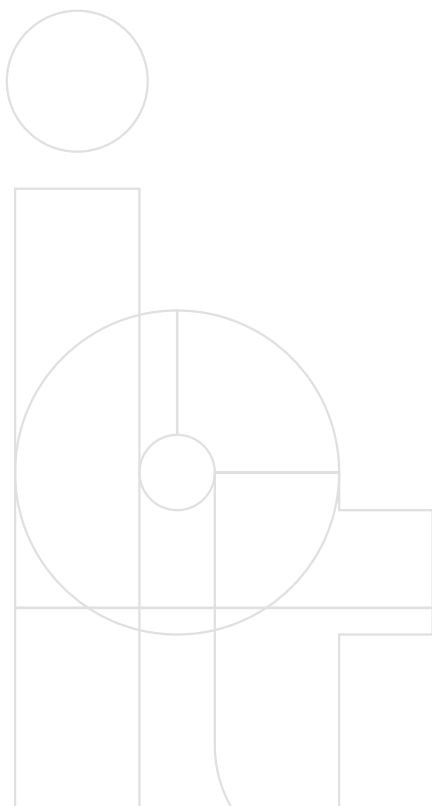
For example, the semiconductor temperature sensor mentioned earlier can provide a reading of 1 microamp per degree Kelvin for its environment. However, does a reading of 273 microamps actually mean that the temperature is *exactly* 273 degrees Kelvin (0 degrees C)? Or could the reading be incorrect to a certain amount of error? Without calibration, it is difficult to be completely certain, although it is a good estimate of the temperature.

Other temperature sensors are even more problematic. For example, the Resistance Temperature Detectors can be a hundred times more accurate than a semiconductor temperature sensor, but without correction, its readings are quite useless. The RTD requires precise signal conditioning, linearization, and calibration to achieve that accuracy.

These corrections often are applied digitally, as when a reading from the RTD is first converted to a digital value, and then the correction is applied. Indeed, different types of RTDs need different types of corrections. For example, a platinum RTD has two distinct relationships to temperature with different polynomial equations describing its resistance above and below 0 degrees C.

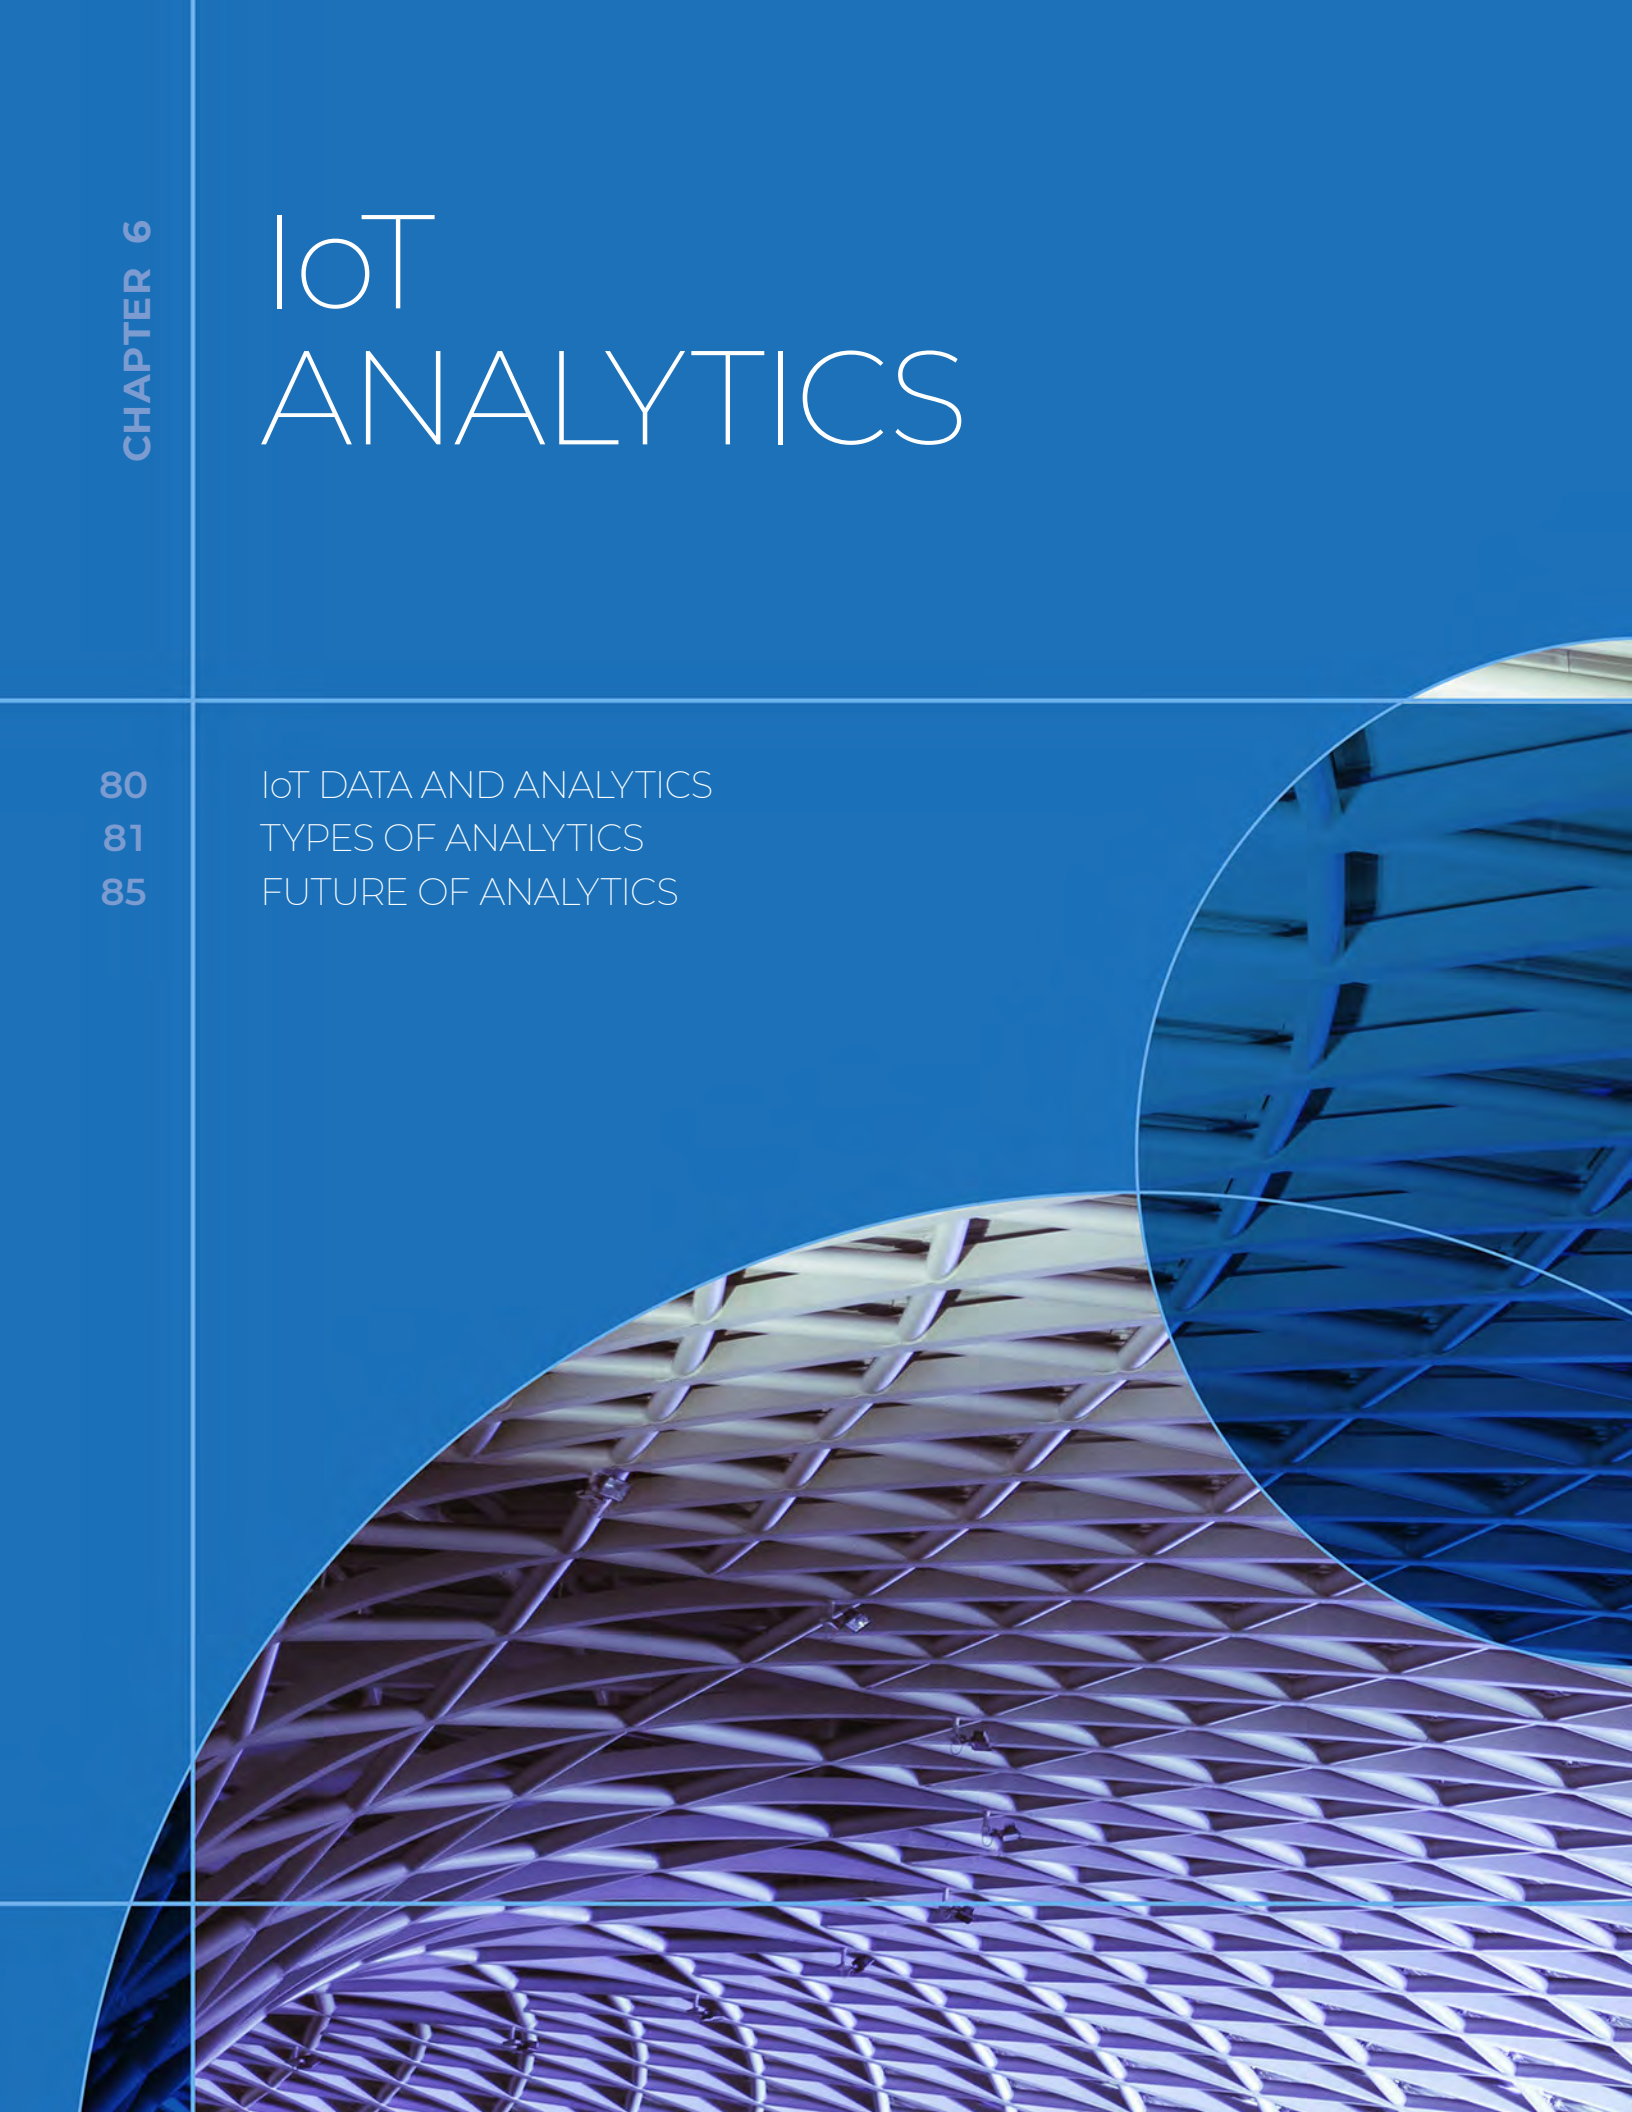
In an RTD, to achieve the best accuracy, the measured signal can be corrected using a variety of techniques, such as direct math, single linear approximation, or piecewise linear approximation. Each has its advantages and disadvantages.

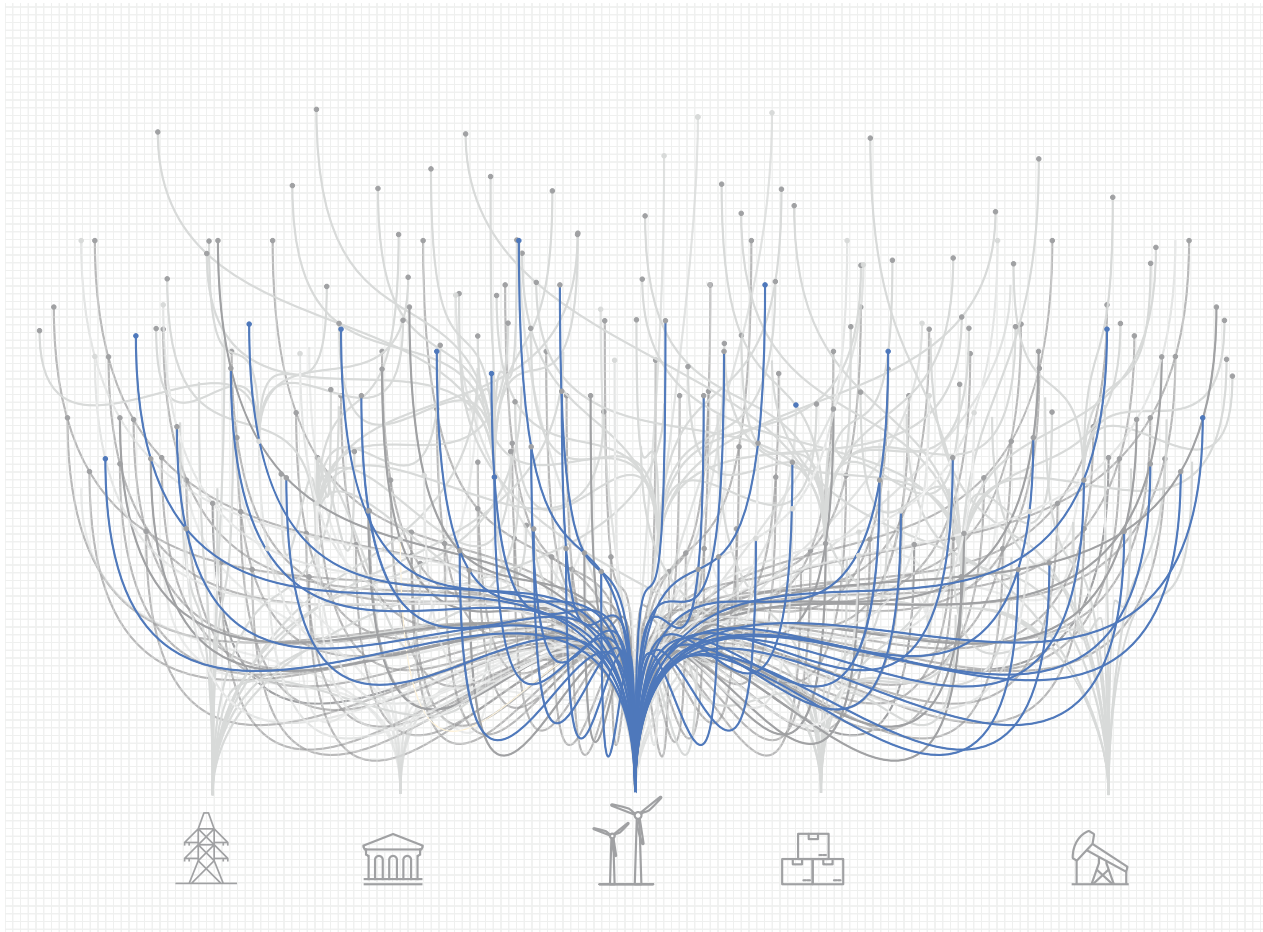
It is beyond the scope of this book to describe how to correct the readings from sensors (for example, for correcting the RTD measurement). Suffice it to say that developers designing IoT applications must take linearization and calibration into account for the specific needs of their application—particularly if the desired accuracy is important to its function.



IoT ANALYTICS

80	IoT DATA AND ANALYTICS
81	TYPES OF ANALYTICS
85	FUTURE OF ANALYTICS





IoT ANALYTICS

In the context of IoT applications, looking at received data and finding meaningful patterns in that data is the basis of analytics. These patterns could describe the state of the data, predict an outcome, find correlations between variables, project trends in the data, and a lot more. Analytics are used in many aspects of business, from marketing to risk management. In this chapter, we will discuss analytics as it relates to IoT data.

Over time, IoT applications can generate vast amounts of data, both because of the large number of units expected in the future, as well as transmissions over years of service. This is part of the Big Data revolution that is much hyped in the media. For example, the Aeris IoT network manages traffic with more than one billion IoT events each day. The more IoT data points being collected, the more need for sophisticated analytics to understand and gain value from the patterns.

New ways to process and store computing data has made it possible to apply analytics to business problems faster and at a greater scale than ever before. Successful organizations take advantage of these tools and analyze the data from IoT deployments so as to gain insights into everything from how to streamline manufacturing processes to the satisfaction levels of its customers.

IoT DATA AND ANALYTICS

IoT devices usually report data in constant streams or periodic messages, and these must be processed both in real time for immediate decisions and alerts, as well as in batches for deeper insights into learning patterns and behaviors.

Learning models built from batch and streaming analytics are used to define thresholds for real-time analytics. This is a dramatic shift from traditional analytics methods that mainly were single-file-oriented data processing programs. Today, real-time analytics are possible, and business-necessary, on streaming IoT data.

Another key to IoT analytics processing has been the development of new tools, open-source distributed storage, and distributed processing frameworks. The Hadoop ecosystem (Hadoop Distributed File System), as well as products such as Riak, Cassandra, MongoDB, Apache HBase, CouchDB, Redis, etc., are being leveraged for Big Data storage and analysis. This allows processing of very large and streaming data sets over computer clusters. Hadoop and components can be deployed as a cloud-computing service by smaller organizations.

With these scalable technologies capable of analyzing and storing streams of Big Data, businesses can use various types of analytics to better understand their collected data from IoT applications and devices.

TYPES OF ANALYTICS

Analytics can be grouped into four broad categories: descriptive analytics, diagnostic analytics, predictive analytics, and prescriptive analytics.

Descriptive Analytics

Descriptive analytics, also called descriptive statistics, provide a numerical or graphical representation of the data that is available right now. It provides a way to express, in absolute, unambiguous terms, a quantitative measurement of the current state. This analysis can draw conclusions from the past as well.

In a broad sense, descriptive analysis answers these questions:

- What happened?
- How often did it happen?
- How reliable was it?
- How accurate was it?

Knowing the current status of IoT data provides a baseline against which to compare future states. It is possible to compare basic data from the past to the present, tracking progress along the way. Descriptive analytical tools can be as simple as tracking website traffic or more complicated, such as cluster analysis used in market research.

Diagnostic Analytics

This type of analytics often is merged with descriptive analytics, and, together, they can give data greater interactivity. Where descriptive analytics asks “what happened?”, diagnostic analytics asks “why did this happen?” The diagnostic tools can be applied to the data to look for the root causes behind the results observed in the original data.

Usage-based insurance (UBI) implemented with vehicle telematics is one example of descriptive and diagnostic analytics in action, in tandem. This type of vehicle insurance establishes the driver’s insurance premium rates on behavior that is tracked via a GPS-enabled cellular transmitter in the vehicle. The distance a person drives, when a vehicle is driven, and where the vehicle is tracked, as well as other attributes, are used to calculate the insurance cost.



By finding patterns and trends in the data, it may be possible to predict future results.

Predictive Analytics

Prediction is one of the main reasons that businesses use analytics in the first place: predictive analytics provide a means of projecting what will happen next, based on what has happened in the past. By finding patterns and trends in the data, it may be possible to predict future results.

Of course, assuming that future behavior will be the same as past behavior isn't always the correct call. Although, unlike the stock market or consumer purchasing habits, machine behavior generally is highly predictable. In a factory, vibration and temperature data broadcast from an IoT-connected device can indicate, with a high degree of accuracy, when a machine needs preventive maintenance.

Businesses can use predictive analysis in their own IoT deployments as part of supply chain management and manufacturing processes to increase efficiencies.

For example, brake balancing in trucking fleets is a complex and expensive maintenance issue. Without regular maintenance, the risk of a truck jackknifing on the highway is high. It takes a highly trained technician significant time to check the combination of brake temperature and pressure to know when to make an adjustment to the vehicle's brakes.

In Michael Lawrence-Smith's study, "Cooperating Artificial Neural and Knowledge-Based Systems in a Truck Fleet Brake-Balance Application," he describes how machine learning techniques used predictive analysis to improve brake maintenance. These computer-aided systems have a 90% success rate at predicting when to replace brakes, resulting in an annual savings of at least \$100,000 for larger trucking companies.



Prescriptive Analytics

Prescriptive analytics is the logical next step from predictive analytics. It asks what a business should do based on the data that has been collected and analyzed. Prescriptive analytics uses models to both recommend actions and forecast outcomes in order to reduce risk.

Just as descriptive and diagnostic analytics work well together, predictive and prescriptive analytics also work hand-in-hand. As past data is used to calculate future results, prescriptive analytics can be used to make better choices and take advantage of opportunities.

Google's self-driving cars, for example, use prescriptive analytics to make countless driving decisions. The cars communicate with the cloud using IoT systems to obtain data on traffic and weather, which becomes part of its driving computations. The vehicle's on-board computers apply machine learning to the problem of what a car should do based upon predictions of future outcomes. For example, the car's computer may predict traffic based on the time of day and then determine what route to take and what speed to travel safely.

But analytics only will predict, for example, that a machine part will break if it has seen signatures of such failures before. Therefore, gathering data from large deployments is very important.



FUTURE OF ANALYTICS


Organizations continue to expect more from their analytics tools and investments. The use of Artificial Intelligence (AI) techniques and algorithmic approaches, such as Machine Learning (ML) and Deep Learning (DL), provide greater insights into the data than ever before possible with traditional batch-oriented data crunching. Indeed, as the number of IoT devices grow rapidly over the next few years, the ability to provide information processing at scale becomes critical.



If the analysis cannot be performed in the required intervals for the resulting actions to be valid, the quality of business decision outcomes drops. Companies must look at new architectures and solutions, including micro-services and APIs to external processing systems, for example, to interact with their existing mechanisms for analysis, and adapt accordingly.

The need for speed—including real-time or near real-time analytics—is increasingly important to act on data in motion. These analytics systems must ingest the streaming IoT data seamlessly, analyze it in the context of past data and history stored in Big Data systems, and act, often automatically without human intervention, when and where appropriate, at the scale needed. Often, new hardware architectures will reduce the processing constraints encountered with older computing systems. For example, using high performance video graphics processing units (GPU), with their hundreds to thousands of computing cores, can provide processing gains that easily exceed traditional server capabilities.

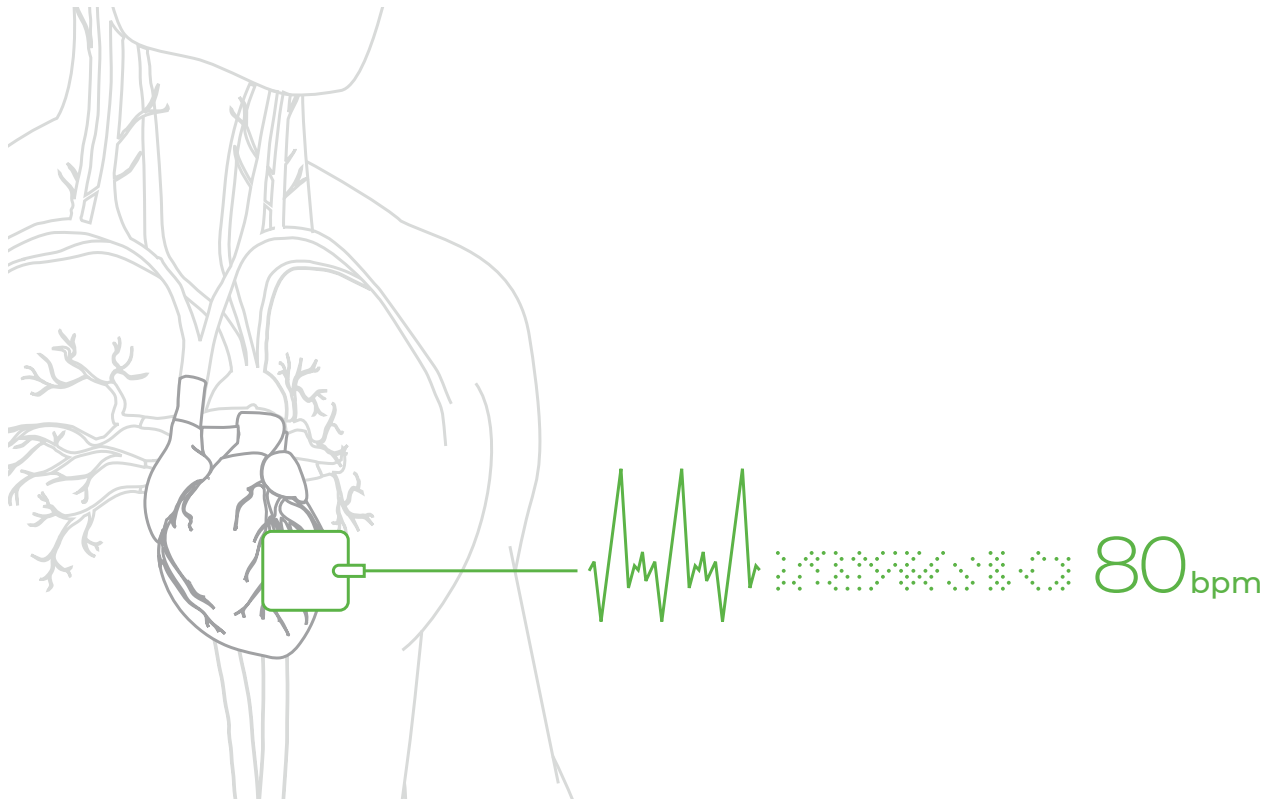
Finally, the ability to deploy analytics platforms in the cloud plays a major role in this space. IoT data is being generated globally, and local cloud services allow the transport of that data to reach the data processing systems without overwhelming long-range networks or crossing national boundaries that may be regulated.



SCHEDULING, ENCODING, AND PROCESSING

89	DATA TRANSMISSION SCHEDULES
91	UDP OR TCP
93	CONTENT ENCODING / TRANSPORT PROTOCOLS
98	GATEWAYS
98	APPLICATION SERVERS
99	CLOUD COMPUTING
100	FOG COMPUTING





SCHEDULING, ENCODING, AND PROCESSING

Data and sensor readings generally are transmitted to Internet of Things and Machine-to-Machine application programs for processing, storage, and business actions.

This may be a relatively short-range transmission from the IoT devices. Sensor readings can be delivered to a smartphone application using a short-range wireless technology, such as Bluetooth, ZigBee, or Wi-Fi, for an action by the owner of the smartphone. For example, a heart rate monitor may send heartbeats-per-minute to a smartphone application during exercise, and this can be monitored to modify the specific physical activity. The data can be logged by the application, possibly into a cloud application, to ensure that the desired fitness goals are being met.

For other IoT applications, the data may be sent over a longer-range transport to cloud systems, application servers, and programs where it is processed for actions or stored for analytics. The data (or patterns in the data) may lead to business actions, including automatic functions performed programmatically, if appropriate for that specific application. For example, an airbag deployment notification from a vehicle can be sent to an automotive Telematics Service Provider (TSP) that contacts the driver and connects them to public safety personnel for dispatch of emergency services.

This chapter describes the systems and methods used to encode, transmit, store, and process the data in a server application.

DATA TRANSMISSION SCHEDULES

Devices may transmit their data in real time, a scheduled rate, or when the device firmware requests a report of an event. Application servers also can initiate a transmission from a device by polling it with an appropriate control message sent to the device.

Devices that send their data continuously in real time or near real time are “streaming” applications. The processing of this streamed data requires systems capable of handling the high throughput from a large number of devices, particularly if the content is to be analyzed in real time for specific actions at a remote site.

The cost of transmitting real-time streaming data on “metered” communications networks that charge for “quantity of bytes sent” may be prohibitive for many applications.

For a large deployment, the requirements for processing the data may be large enough to clearly require a commercial cloud service provider that can handle the necessary throughput with bandwidth, performance, and high-availability systems to deal with the data stream. The application may need to store the data for long periods of time for long-term analytics functions to be useful.

Scheduled Transmissions

In some applications, devices transmit on a regular schedule—sometimes sleeping to conserve power until they are woken up by scheduled timers or to report an unscheduled event.

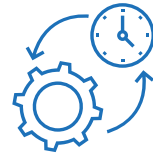
Devices with accurate time (such as those equipped with GPS receivers) must be careful when using regularly scheduled transmissions. In large deployments, if all devices were to wake and transmit at exactly the same time, the simultaneous connection attempts could overwhelm the connectivity paths and the server systems that receive and process the data. If possible, randomizing the time of the transmissions can have a major positive impact on the capacity requirements of the connectivity and the server systems.

There are simple ways to achieve this randomization. For example, a device identification number—such as the last four digits of the Mobile Directory Number (MDN) or the Mobile Station ISDN (MSISDN) in a cellular device (modulo 3600 to bring it into the correct range)—can be used to select the “number of seconds past the hour” when a regular transmission is sent.

Transmit On-Demand

In most IoT applications, it is typical for the device to transmit “on demand” when an event requires it to do so. For example, a business or residential security system can transmit a signal when a break-in occurs; a car may transmit an accident notification when an airbag deploys; or when the driver pushes a concierge button for assistance. These generally are sporadic enough and temporally spaced that they do not create traffic (and server system) spikes.

Often, devices that transmit to report sporadic events also are set to transmit a periodic, regular “heartbeat” to report their condition and health. These heartbeat transmissions also should be randomized.



UDP OR TCP

We often are asked whether a device should use User Datagram Protocol (UDP) packets or Transmission Control Protocol (TCP) streaming sessions for the data. The answer, not surprisingly, is: “It depends!”

The Internet Engineering Task Force (IETF) has detailed definitions for these two protocols, but let’s briefly describe them to understand why one may be better than the other for certain IoT data transmissions.

It is important to note that both UDP and TCP are used over an underlying Internet Protocol (IP) data connection.

User Datagram Protocol (UDP)

The UDP format was first defined in an IETF Request for Comment (RFC) specification—RFC 768. This protocol provides a procedure for application programs to send messages to other programs with a minimum of protocol overhead. This protocol is transaction-oriented, but delivery and duplicate protection are not guaranteed.


If an application requires ordered, reliable delivery of streams of data, UDP is not the preferred protocol. However, the UDP format has lower overhead than TCP—i.e., fewer bytes are sent in the headers of the packets in UDP than TCP.

Transmission Control Protocol (TCP)

The TCP format was first defined in an IETF RFC specification—RFC 761. TCP is a connection-oriented, end-to-end reliable protocol that is intended for use as a highly reliable host-to-host protocol in IP networks and especially in interconnected systems of such networks.

TCP requires that a connection be opened and managed for the duration of the data transmission on an IP network. Within the protocol, transmitted and received packets are acknowledged by the device and the servers.

This format has more overhead than UDP—i.e., more bytes are sent in the headers of the packets in TCP than UDP.



Which to Use?

In general, the choice of UDP vs. TCP must take into account:

- The desired balance between the reliability of TCP and the lower cost of UDP, since UDP uses fewer bytes of overhead to transmit the same amount of application data.
- The increased complexity of TCP, where the module must open a data stream to a remote server where programs await connections.
- Careful design of TCP server programs to allow easy scaling as the number of deployed devices increases for an IoT application.
- A requirement for the acknowledgments provided by TCP sessions.
- UDP is suited for real-time applications that can tolerate packet loss. TCP is suited for applications that can tolerate delay but not packet loss.

However, it also is important to note that using these two protocols is not mutually exclusive for a given IoT application.

For some data, a simple transmission of a UDP packet to a remote server may be quite sufficient—including possibly using independent acknowledgments also via UDP. If an acknowledgment is expected, but not received, either side can retry intelligently (i.e., with limits on the number of retries, variable delays between retries, etc.).



For other data, even in the same IoT application perhaps, a device may open a TCP connection to a server and communicate with the higher reliability of a TCP streaming session to a program that accepts these connections and transmissions, while providing direct real-time acknowledgements.

Often, the amount of data for a particular data set may *require* TCP. For example, if a device needs to transmit a large amount of data for a particular set of gathered information (i.e., more than a kilobyte), it generally is better to use TCP since the consequences of an error during transmission via UDP could mean that the entire data set might require a complete retransmission.

It should be noted that in modern IP communications—including cellular IP data—this unreliability concern is low, and UDP should suffice for a significant set of the data transmitted for a particular IoT application.

CONTENT ENCODING / TRANSPORT PROTOCOLS

When a device transmits its data to the servers and receives commands and instructions from the servers, an encoding format is required for the information sent in both directions. In every application, the devices and servers must formally agree on the format and information that is transmitted.

Proprietary Format

Devices and servers for a particular IoT application could choose to use a proprietary format for the data encoding. This allows the devices and servers to encode, decode, and interpret the content in ways unique to the needs of that application. This often can minimize the amount of data sent in any connection session.

Proprietary formats are more difficult to implement initially—since they must be quite complete for that application to be deployed—as well as difficult to maintain and update later when changes are needed. Most proprietary formats are not easily extensible.

Common Industry Formats for IoT

In addition to proprietary formats and early standardized formats, such as eXtended Markup Language (XML), there are some industry formats and protocols in use for IoT data communications for messaging needs:

- JavaScript Object Notation (JSON)
- Constrained Application Protocol (CoAP)
- Message Queuing Telemetry Transport (MQTT)
- Extensible Messaging and Presence Protocol (XMPP)

These protocols fall into two basic categories: human-readable (JSON, XMPP) and non-human-readable (CoAP, MQTT).

The human-readable ones generally are much more verbose, but far easier to debug during application development and subsequent operations. The other, non-human-readable ones are lighter weight and efficient and can minimize the amount of data sent over the communications path.

Each format (and protocol) has pros and cons when used for IoT. The choice depends on the needs of the application, the bandwidth of the communications network, the compute power in the sensor or remote device, and other factors.

The choice depends on the needs of the application, the bandwidth of the network, the compute power, as well as other factors.

JSON

JSON is an open-standard format that sends key-value pairs of information. The “key” generally is the attribute or description of the content sent in the “value”. The protocol is described in RFC 7159 from the IETF.¹

The JSON format is human-readable and language independent, and public code for parsing and generating JSON text data is readily available in a variety of programming languages. The format is effectively self-describing since the definition and value are right next to each other.

For example, the following simplified text illustrates the encoding of a temperature reading of 25 degrees Centigrade from a sensor with a hypothetical sensorID of 123456789:

```
{  
  "sensorID" : "123456789" ,  
  "temperature" : "25" ,  
  "units" : "Centigrade"  
}
```

As you can see, the JSON content is verbose and very human-readable. The key-value pairs immediately identify the attribute and its value—picking appropriate terminology for the keys that are meaningful is, of course, important for this capability to be useful.

JSON format messages also can be extended readily. For example, the physical location and manufacturer might be added, along with a time stamp noting the time that the temperature was measured.

Of course, the presence of this additional information depends on whether it should be transmitted. In the above example, the value of sensorID could be used to look up the physical location in a server database (assuming it was stored there at installation of the sensor). Moreover, sending a time stamp from the device for each transmission can be very useful since it provides the time when the data was collected (assuming the device knows that time information, of course).

¹ See www.ietf.org/rfc/rfc7159.txt for more information.

CoAP

As the name implies, CoAP is a format and protocol intended for use in bandwidth-limited networks or where minimizing the size of each message transmission is important. The core of the protocol is described in RFC 7252 from the IETF, although extensions to add unique requirements for IoT applications currently are in development.

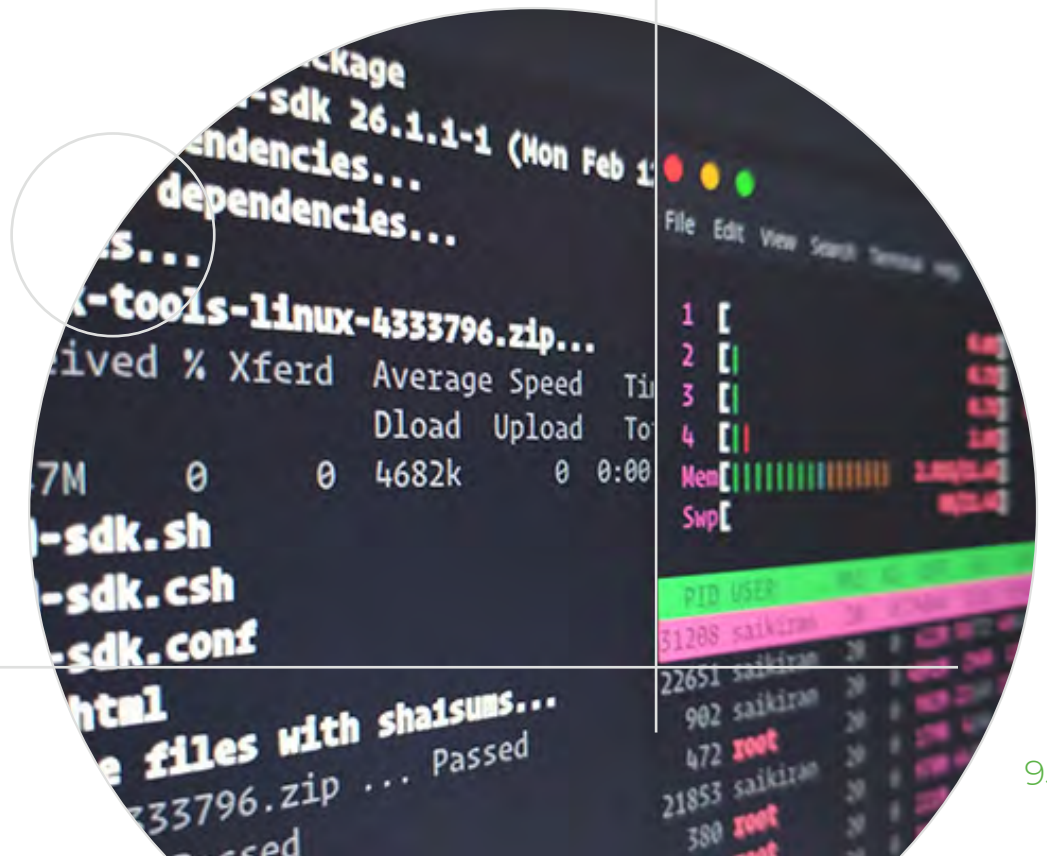
CoAP is a simple protocol that is well-suited for transmissions from small electronic components, such as sensors, and also can be used to control the devices from remote servers. CoAP includes the concept of “multi-cast” (or “one to many”) group communication, where many devices can receive the control information at the same time.

The protocol provides two types of messages: requests and responses using a “type-length-value” (TLV) coding that is different from the JSON format. CoAP messages are sent using a UDP transport to adhere to the concept of low overhead for the messages.

MQTT

MQTT is another light-weight messaging protocol that is designed for data transmissions from devices operating in bandwidth-limited networks. The devices transmit the data to message brokers that then are responsible for sending the content of the messages to clients who are interested in that data and who subscribe to the data feed.

This mechanism is the essence of a “publish-subscribe” approach, where data from a device is published to a broker, and subscribers to that broker can access the data.



Originally developed by IBM, the MQTT protocol was transferred to the OASIS¹ standards body and now is supported by that entity.

MQTT originally was designed for the IoT markets for devices transmitting using TCP / IP. To allow simpler electronic devices (such as sensors) to use this protocol, a version called MQTT for Sensor Networks (MQTT-SN) has been released to extend the protocol beyond TCP / IP.

XMPP

XMPP is an open-standard communications protocol for messages based on XML. It is intended for near real-time exchange of messages between two (or more) elements on any network. Like XML, it is extensible and also can be used for publish-subscribe message systems.

There are multiple RFCs from the IETF that specify the XMPP standards: the core ones are RFC 3922, 3923, 6120, 6121, and 7622, although the XMPP Standards Foundation² is actively extending XMPP further.

The XMPP protocol evolved from an earlier open-standard protocol called Jabber and was used for Instant Messaging (IM) services, as well as Voice over IP (VoIP) control messages. In this last application, XMPP competes with the Session Initiation Protocol (SIP).

When XMPP extensions are used for publish-subscribe services, they are useful for IoT data applications. However, like JSON, they are human-readable and verbose—perhaps even more verbose than JSON due to the XML roots. This may make it difficult for a small sensor to encode XMPP directly, but a communications device could make the necessary conversion from raw sensor data.

In XMPP, binary files and content can be encoded (using base64 conversion of the binary data to text) and sent using XMPP, but this is likely to use more overhead than is desirable for IoT applications.



The gateway is a good location in the communications path to implement the data encoding, as well as security best practices.

GATEWAYS

In most low-cost sensors—even the newer ones that “speak IP”—it is difficult to provide the data encoding and decoding functions within the sensor. Often, these sensing devices use short-range communication paths—either wireless or wired—to a device with more computing capacity that actually encodes the data and transmits to a remote server.

This device may be a physical unit serving a single sensor and associated application. More often, it is a gateway—a product with multiple short-range wireless and wired connections to local sensors and a long-range wireless or wired connection to the remote IoT servers.

For example, gateways used in home automation applications typically communicate with sensors using Bluetooth, ZigBee, and low-power Wi-Fi, and to the remote servers with cellular or wired Ethernet IP connections.

The gateway is a good location in the communications path to implement the data encoding, as well as security best practices, with software agents that take the raw information from the sensors and encode the data in the formats described above. The gateway also could implement encryption algorithms to protect the data.

For large-scale deployments, the application servers literally must be running continuously with high availability and processing redundancy.

APPLICATION SERVERS

Remote data is transmitted to application programs running on the servers that may be dedicated to the task of processing that data—whether it is streaming data or message oriented.

Typically, these servers are deployed in data centers on the customer premises or in standalone data centers. The programs on the servers receive the data and process them for the specific business action of the IoT application. This may include storing the data in traditional databases, filtering for erroneous information, alerting when the information is outside pre-determined bounds, displaying the data or reports, etc. The needs vary greatly.



Often, remote devices, even those that are transmitting lightly, cannot tolerate server downtime for any significant duration. Thus, processes and network infrastructure to automatically balance the loads on redundant servers, including at multiple sites, are critical.

For large-scale deployments, the application servers literally must be running continuously with high availability and processing redundancy (including geographic redundancy), particularly for mission-critical applications. With the projected growth of the IoT market, this will place an immense burden on servers and data centers. This need creates a significant capital and operations cost of systems, physical site maintenance, power distribution, cooling, and more.

The choice of which server platforms, operating systems, programming languages, etc., is dependent entirely on the entities deploying the IoT application. Traditional IT departments generally have all the relevant expertise to make these decisions for the companies.

In most cases, however, where massive growth is expected to occur, IoT deployments should consider taking advantage of newer IT deployment architectures, like “cloud computing”, and data traffic reduction methods, such as “fog computing”.

CLOUD COMPUTING

In recent years, the phrase “cloud computing” or simply “the cloud” has been coined to describe the systems that allow processing and storage of data in extremely large data centers for a fee. Cloud vendors provide the ability and flexibility to start and stop computing and storage of data, while providing the networking resources based on the specific needs of the customers and their applications using these cloud services.

This has transferred the need for entities and corporations to maintain their own physical hardware systems, data centers, and data networks, etc., to the cloud vendors. This eliminates the “traditional” operational burdens of physical site maintenance, electrical power management, environmental conditioning, and system redundancy.



The specific compute, storage, and transport requirements for the cloud customers then can be adjusted fairly dynamically to conform to the needs of the applications being executed. The latest techniques and software for managing large amounts of data can be applied to the data gathered from the devices in the IoT applications.

These compute elements, storage, and data transport are, of course, provided for a fee. The charges vary but often can be quite high for large-scale IoT applications and large numbers of device deployments.

FOG COMPUTING

The volume of data gathered from a large number of sensors and devices could overwhelm the IoT data communications path (transmission and connectivity) or the remote storage capacity and server systems that process the data at customer sites.

While cloud solutions do alleviate this problem, the cost could be very expensive—particularly for streaming applications. Often, a general approach to remote data gathering is a “transmit everything and process in the cloud” implementation.

However, if actions based on the data must be processed in real time or near real time, it may be better to process or filter the data remotely—at the device, or elsewhere hierarchically in the data path—before it gets to the servers. This processing and filtering has been termed “fog computing” by Cisco.

Fog computing is not without its issues and concerns. If the filtering removes essential information that could be better processed at a central site (such as the cloud) to determine patterns, its use could result in a weaker application.


Sometimes, the specific filters used at the remote device may need to evolve and change. Thus, IoT devices must be programmable, or sufficiently configurable, to change the specific data that is transmitted, thereby increasing the complexity of the overall solution.



One significant advantage of fog computing is the concern about IoT security. Good security practices can be implemented farther away from the central servers, where a device (or groups of devices) that have been compromised could be isolated, perhaps limiting damage to the overall application deployment.

It also reduces the transport costs of sending a lot of data—much of which may be meaningless, repetitive, or simply not needed—on metered transports where the transport of large sets of data can be expensive.

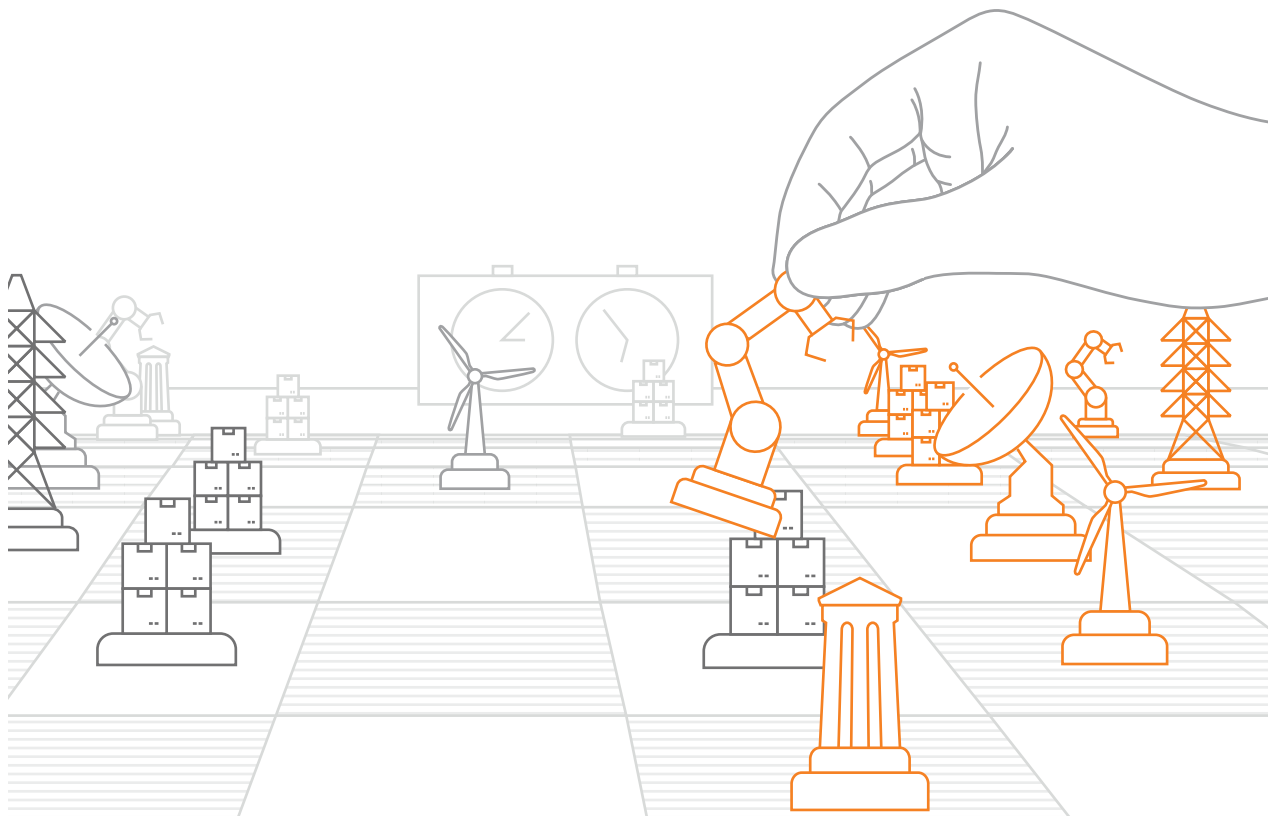
-
-
-



IMPLEMENTING AN IoT SOLUTION

104	SUPPLY CHAIN MANAGEMENT
104	CELLULAR OPERATOR SELECTION
106	CLOUD SYSTEM SELECTION
107	PLATFORM SELECTION
107	NETWORK OPERATOR SERVICE LEVEL AGREEMENT
108	DEVICE CERTIFICATION
109	NORMAL OPERATION CONSIDERATIONS
111	APPLICATION COMMUNICATIONS CALL FLOW
112	CUSTOMER SUPPORT PROCESS





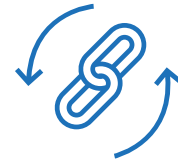
IMPLEMENTING AN IoT SOLUTION

An Internet of Things application deployment has to either increase business revenue or reduce business costs (or both), otherwise there's no reason for a company to pursue it. Either of these objectives can provide a return on investment. It's up to the product manager to determine the specific goals and measurements of this ROI.

Related drivers for IoT projects can be new regulations and industry requirements, greater efficiencies, more consistent control over processes, visibility into patterns or opportunities, or gaining competitive features that can meet customer requirements. As you build your IoT business model, these factors will weigh differently depending on product needs and the industry.

SUPPLY CHAIN MANAGEMENT

Supply chain management refers to planning for the flow of materials and services into and out of the business, and managing all the goods required to make your IoT deployment happen. If your company is building its own IoT devices from scratch, you'll have many materials, parts, and suppliers to account for. If you're assembling devices from ready-made components, you can reduce the number of suppliers. However, even if you buy a complete, off-the-shelf device, it still requires sourcing, testing, and managing supply and demand.



CELLULAR OPERATOR SELECTION

The service provider must be able to deliver several essential requirements for the project, including reliable network connectivity, robust service agreements, effective application integration, cost management tools, and flexible rate plans. If they can't deliver on these prerequisites, they are not going to be the right partner.

To help you select the ultimate service provider with the capacity to manage a successful deployment, you may want to ask these questions of any cellular carrier during the selection process:

- **What are the costs for the entire device lifecycle, not just per kilobyte rates?** Make sure you won't be hit with hidden costs from your cellular operator that drive up your IoT service bill.
- **Can the service provider expand cellular coverage beyond its own cell towers?** Traditional operators only optimize their cellular coverage based on their cost of delivery, and they always prefer to use their own towers, even if the coverage they provide is weak or intermittent. A carrier-agnostic provider, like Aeris, can expand coverage where needed, and will offer the strongest signal, regardless of operator, with no interruption in service.



- **Do they offer remote troubleshooting, as well as hands-on support?** Cellular carriers with remote, real-time troubleshooting capabilities can save you significant costs. Also, an operator with a network operations center support team that deals only with IoT-related issues is going to be more knowledgeable about your devices and connectivity issues.
- **Do they offer a dedicated IoT network?** A network dedicated solely to IoT traffic won't experience the delays caused by crowds of consumer handsets. The lower latency of an IoT-dedicated network means you'll be able to rely on mission-critical transmissions to get through the first time.
- **Do they have APIs for easy integration with your existing systems?** Can the cellular operator provide a full suite of free APIs that let you extend the capabilities of your customer-facing applications and back-office solutions, leveraging business applications, such as those from SAP and Oracle? These applications are integral elements of enterprise resource planning-based supply chains and are linked to back-office systems with APIs.
- **Do they offer pay-per-use, as well as per-device billing plans? Can the cellular operator offer rate plans that are flexible enough to meet your needs?** When managing IoT services, it often makes more sense to go with a pay-per-use plan than with a per-device or fixed-data plan. Pay-per-use is most cost effective for lower-usage device profiles. If your devices have higher-usage levels—10 MB or more—a per-device data plan is your best option.
- **Does their system provide cost management tools that automatically notify the company, or take automatic limiting actions, when a device, or a group of devices, are exceeding their cost models?** Devices sometimes malfunction and “run away”, transmitting more often than they should or retry under conditions where it would be better to avoid network attempts. In this case, it is important to automatically notify company personnel to take action, or even set limits where device operation is blocked to avoid uncontrolled costs.

These are some of the top-level concerns your company should consider when choosing a cellular operator. You'll want to partner with a service provider that suits your business needs and can support your IoT project over the long term.



CLOUD SYSTEM SELECTION

Often, the large number of devices that typically are deployed for an IoT application necessitates the use of a commercial cloud service that provides the performance and high-availability capability required by the application.

This requires customers to carefully ascertain whether the cloud provider has the right tools, cost models, and support for their needs. The right questions are not always clear since some aspect of their service or cost model may be appropriate until a certain size threshold is reached. This “scaling challenge” often is one of the most difficult areas of assessing what is possible.

Consider the following:

- **Does the cloud provider have a cost model that matches what the IoT application can bear?** Can the cloud provider assist you in simulating the costs for your application transport, storage, and analysis needs? If the costs do not meet expectations, the return on investment for the application could fall short and lead to an unsuccessful deployment.
- **Does the cloud provider offer data centers in the countries where the devices are deployed?** In some countries, there are regulations that require that data must not cross past national boundaries and the presence of a local data center may be critical to operating within the regulations of that country. Indeed, an absence of a cloud data center in a vital market may preclude the selection of that provider.
- **Does the cloud provider have the tools to support high-availability deployments? Do your software engineers and operations personnel have the expertise to develop and maintain cloud solutions for your IoT application?** Sometimes, the selection of a cloud provider is guided by the available personnel within your company who have experience with that provider. It may be necessary, however, to hire additional resources or use an IoT platform vendor who can guide you to the best possible solution.



PLATFORM SELECTION

Many companies attempt to provide a platform for IoT solutions. This appears to be an area where it is possible to find hundreds of companies purporting to provide “IoT platforms”. In this noisy environment, it is difficult to assess what the capabilities and features of the platform are, let alone how well they would fit for the requirements of your specific IoT deployment.

Given the large variety of possible IoT applications in many different types of markets and businesses, and the large number of available platforms, it is tough to determine the best one for your needs. Yet, it is important to make the best selection as early as possible, since the wrong selection at the early phase of any IoT application deployment could significantly impact and delay the project.

For more information on platforms, please refer to Chapter 2, The Future of Platforms.

The wrong selection at the early phase of any IoT application deployment could significantly impact and delay the project.

NETWORK OPERATOR SERVICE LEVEL AGREEMENT

The Service Level Agreement (SLA) you negotiate with the operator defines the scope of your contract with the operator. This is where your business defines its relationship with its network provider, so it’s important to specify what will keep your IoT deployment running.



Things to consider in your SLA include:

- What are the expectations for connectivity? How reliable is the operator's network historically?
- What are the geographic restrictions of the operator's network, if any? Some carriers may not guarantee service at all towers or all sections of particular metro areas.
- How long will it take the operator's customer support to acknowledge and then take care of a problem?

When entering into an SLA, make sure the agreement is realistic, actionable, measurable, calculated, well-defined, mutually exclusive, and completely exhaustive in covering all aspects of concerned networking services.


DEVICE CERTIFICATION

Devices must be approved or certified to run on the operator's network. For this certification, the focus generally is on testing the cellular behavior of the device.

One example of this might be the behavior of the retry algorithm used by the device if it fails to connect to the application server in your data center. A continuous retry by thousands of devices at the same time could overload the operator's network. Implementing a random back-off algorithm, and testing it prior to certification dictates better device behavior.

Operator certification also provides an opportunity to use the application host server software to perform additional tests that stress the interaction between the device and the server. Unusual scenarios, such as delayed responses from the server (that might be observed during congestion or server scaling), can be used to see if a device handles them gracefully.

In certain markets, such as the healthcare industry, additional regulations for device performance in medical environments and data privacy rules may apply. Additional certification may be required by standards organizations, regulatory agencies (such as the Federal Communications Commission in the U.S.), or even your customers, particularly if there is end-user integration. Each company deploying such IoT applications must determine how to best meet all the regulations that apply to them.



NORMAL OPERATION CONSIDERATIONS

Here are a few of the concerns to be dealt with when IoT devices are deployed:

- **What is the definition of “normal”?** What are the baseline transmission patterns and server performance measurements?
- **What happens if the IoT device can’t connect to the cloud platform?** In addition to having a random back-off retry algorithm, what will the device do with its data? Remember that stale data would be inaccurate when transmitted too late. The device needs to know when to generate an alarm.
- **What should a mobile IoT device do if it loses its cell signal?** The device needs to know when it is appropriate to hold the data in its queue and retry later.

The range of normal operations will vary for each deployment, so you’ll need to set initial parameters for all aspects of the program. Then you can track performance against this baseline moving forward.



The range of normal operations will vary for each deployment, so you'll need to set initial parameters for all aspects of the program. Then you can track performance against this baseline moving forward.

APPLICATION COMMUNICATIONS CALL FLOW

This is where the details of the IoT transmission are agreed upon. Some design issues are:

- Should the device assume there will be a connection when needed or should it be able to queue data for later delivery?
- Will the application “fire and forget” data or will it wait for an acknowledgment? At the network layer, “fire and forget” means to use UDP protocols for transmission.
- The general call flow is to establish a connection, transmit data, wait for acknowledgment, then disconnect. This generally is a TCP protocol implementation.
- Does the data need to be encrypted? That can increase the amount of data being sent.

Your developers will need to outline each aspect of the IoT application’s call flow, accounting for both standard, predictable behaviors and for outliers.

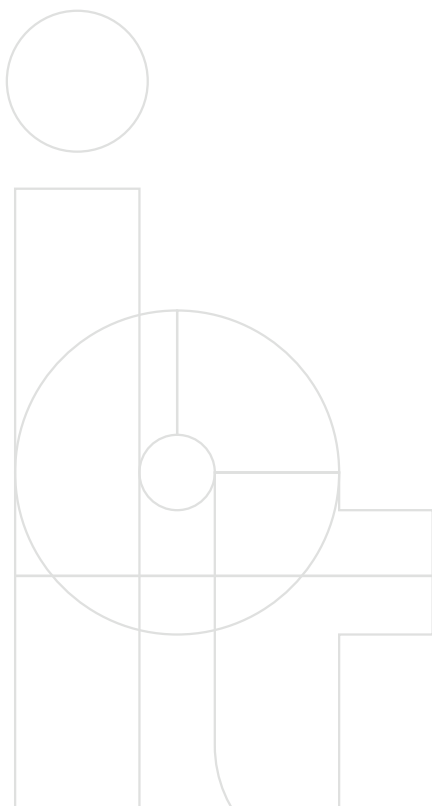


CUSTOMER SUPPORT PROCESS

Support staff will need to be trained on the product features and how to operate them. But it also is very important for the support team to receive training on identifying connectivity issues. This is where a rich set of diagnostic tools from the carrier, if available, become a huge benefit.

If your engineer can log into a portal and see if the device in question has registered on the carrier network and started a data session, then the engineer can observe the recent behavior and immediately can focus the investigation on the root problem. Using this observation, the engineer can provide quick feedback to customers. If these tools are not available, then support sessions are much slower.

Implementing an IoT network project requires a great deal of forethought. But this advance planning pays off in a scalable product with a higher return on investment.



IoT SCALABILITY AND ALTERNATIVE TECHNOLOGIES

117

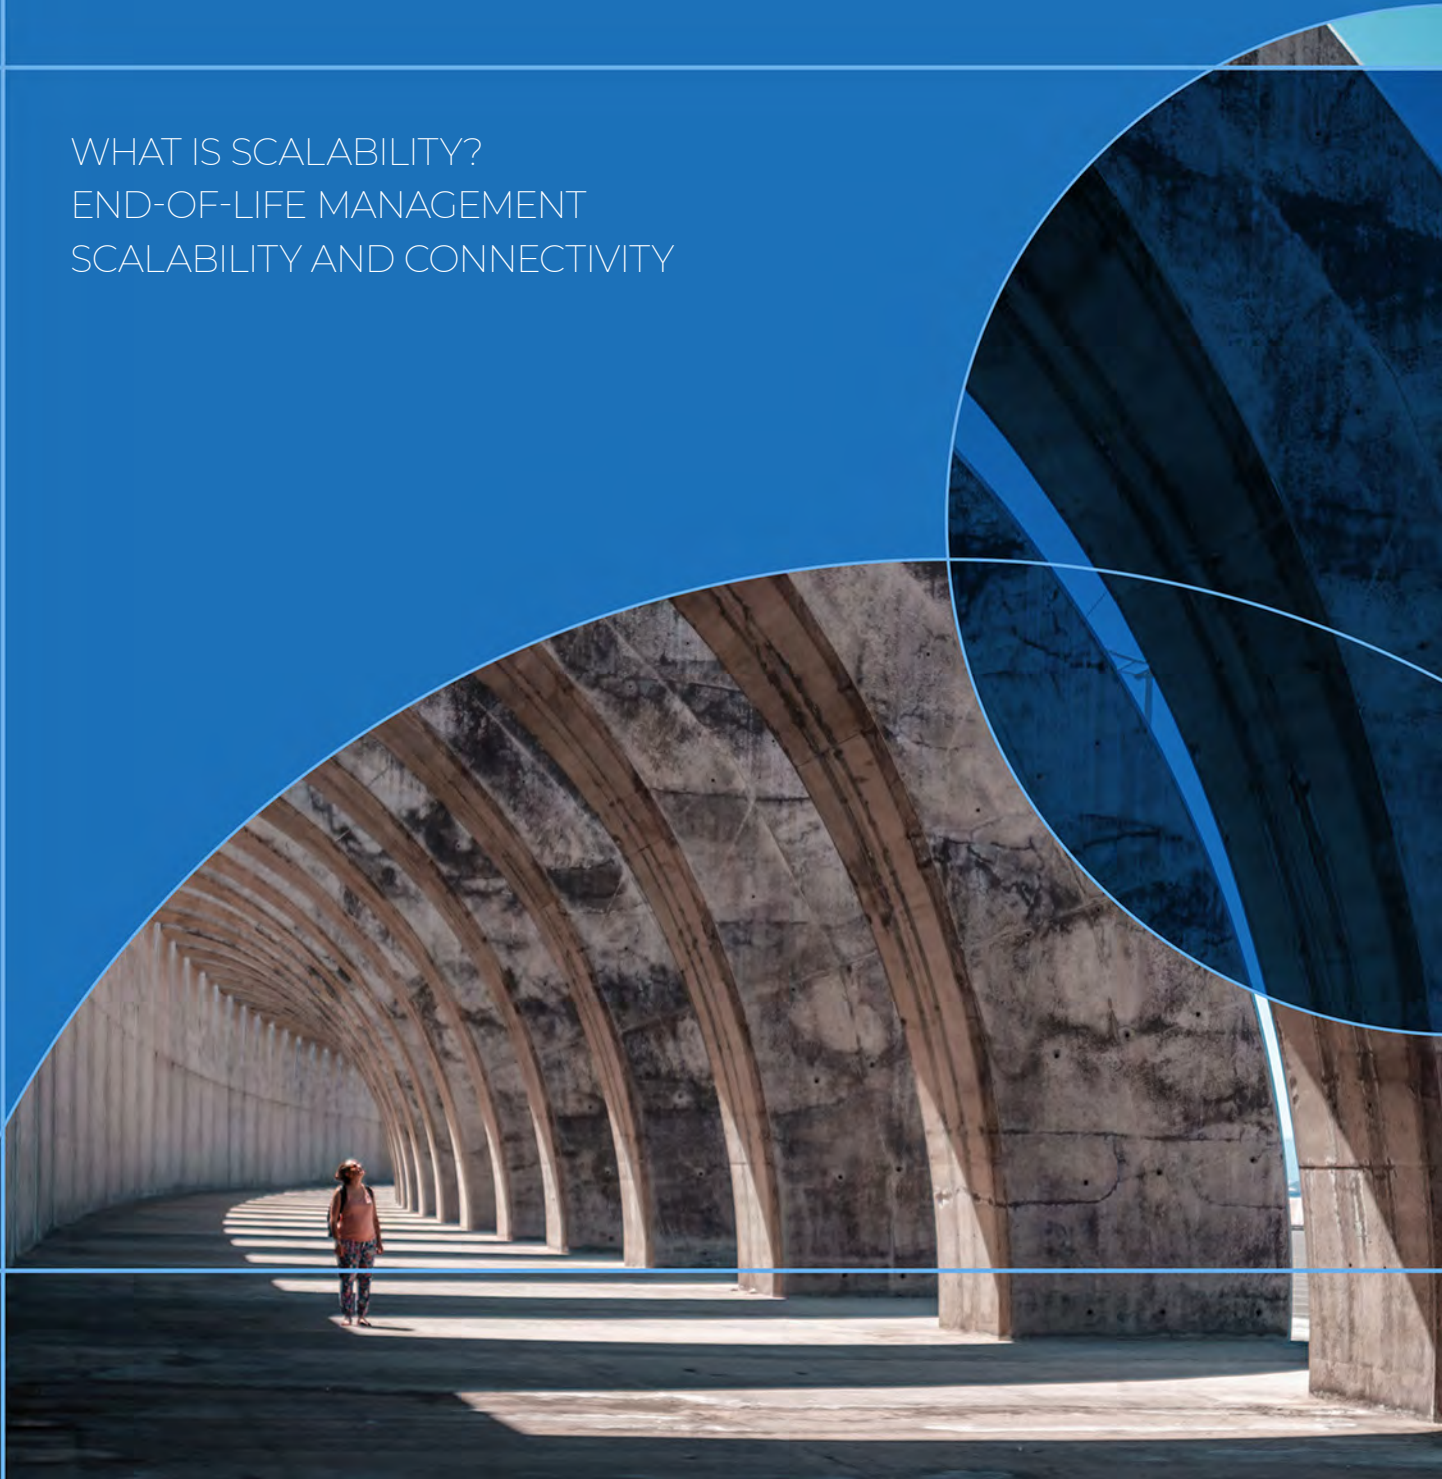
WHAT IS SCALABILITY?

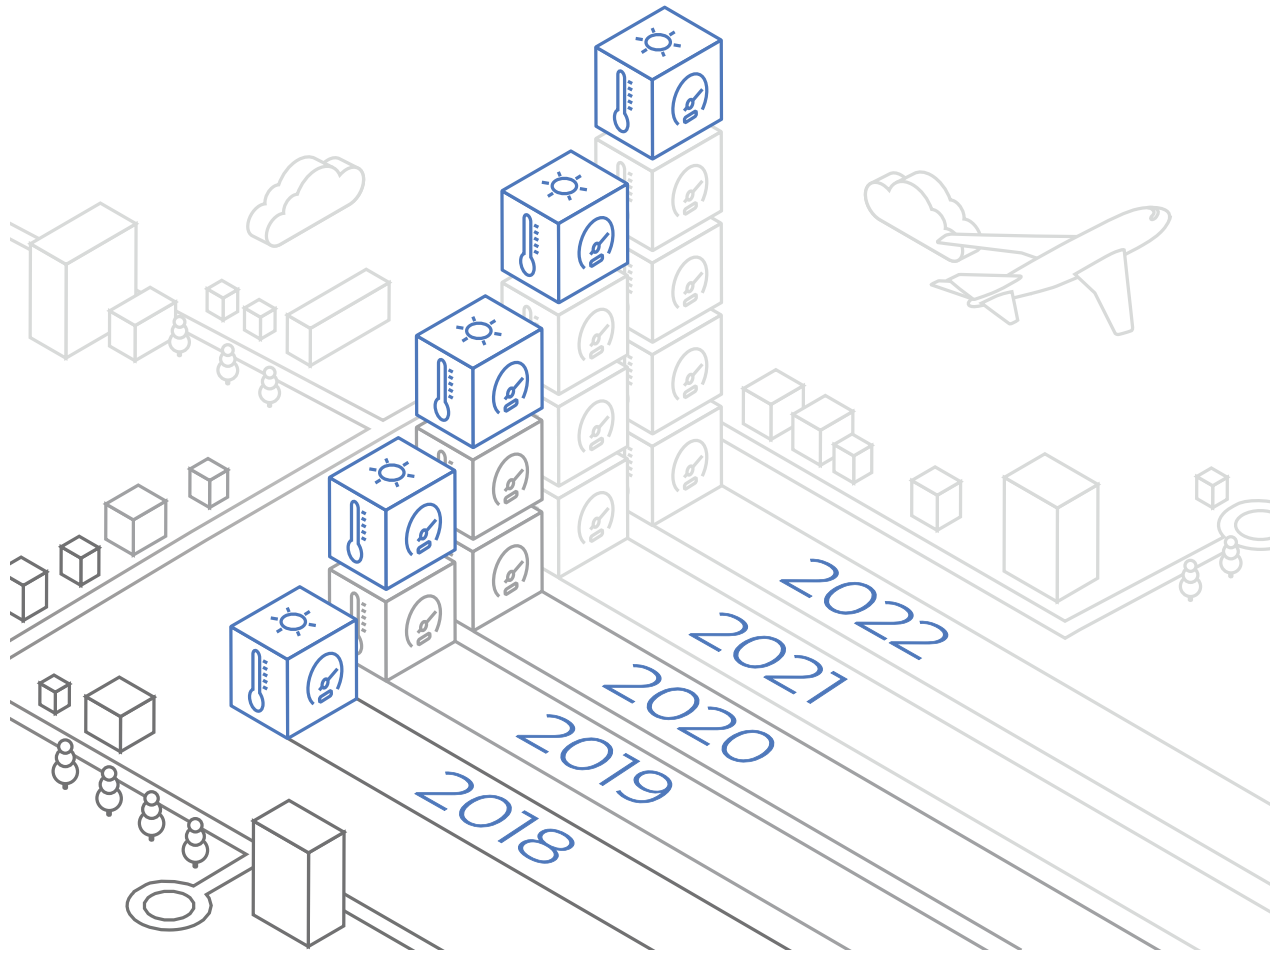
119

END-OF-LIFE MANAGEMENT

120


SCALABILITY AND CONNECTIVITY





IoT SCALABILITY AND ALTERNATIVE TECHNOLOGIES

Over the years, the predictions for growth in the Internet of Things and machine-to-machine markets have been staggering.

1 B connected devices by 2015 ¹	50 B connected devices by 2020 ²	30 B connected devices by 2020 ³
75 B connected devices by 2020 ⁴	31 B connected devices by 2020 ⁵	40.9 B connected devices by 2020 ⁶
20.8 B connected devices by 2020 ⁷	200 B connected devices by 2020 ⁸	

Various predictions of device growth over time

Although the specific predictions and the numbers differ, what is remarkable is that the numbers started extremely high, and only have grown over the years. IoT markets are experiencing explosive growth around the world, and the numbers still are performing at what Gartner calls the “peak of inflated expectations” in its well-known “Hype Cycle” diagrams.

Even if the huge numbers forecasted are inaccurate by large percentages, or even off by a factor of 10 or more, they still represent enormous growth. Indeed, the estimated number of connected devices by 2020 exceeds the projected population of the entire planet by many multiples.

We appear to finally be moving beyond the hype into reality.

¹ “IBM: A World with 1 Trillion Connected Devices,” ReadWrite.com, June 7, 2010.

² “CEO to Shareholders: 50 Billion Connections 2020,” Ericsson.com, April, 13, 2010.

³ “More Than 30 Billion Devices Will Wirelessly Connect to the Internet of Everything in 2020,” ABI Research, May 9, 2013.

⁴ “Morgan Stanley: 75 Billion Devices Will Be Connected to the Internet of Things by 2020,” Business Insider, October 2, 2013.

⁵ “The Internet of Things in 2020,” VisualCapitalist.com, August 23, 2014.

⁶ “The Internet of Things Will Drive Wireless Connected Devices to 40.9 Billion in 2020,” ABI Research, August 20, 2014.

⁷ “Gartner Says 6.4 Billion Connected ‘Things’ Will Be in Use in 2016, Up 30 Percent From 2015,” Gartner, November 10, 2015.

⁸ “A Guide to the Internet of Things Infographic,” Intel.com, 2016.

This explosive growth needs to be managed and planned if we are going to come close to the predictions for what these markets and industries can do for all of us. Although wireline solutions likely will continue to be the dominant technology for overall IoT connections, wireless technologies, including cellular as well as LPWA deployments, will see tremendous growth over the next several decades, driven largely by deployments that require mobility, like automobiles and asset tracking.

All of this anticipated growth in the IoT markets will bring new challenges:

- Scaling for growth in the numbers of devices and applications.
- Providing effective security solutions for the content and solutions (as discussed in the previous chapter).
- Storing the data and providing rapid analysis for action.
- Deploying new wireless and wired connectivity technologies for the increased traffic.
- Managing the connectivity and device “subscriptions” for large numbers of devices.


This chapter will briefly review some alternative transport technologies that likely are to be used for large-scale IoT deployments.

WHAT IS SCALABILITY?

In the context of IoT, scalability is the ability to grow the application, the solution, and the platform to keep up with the projected growth in the number of devices, the data traffic from these devices, the applications servers that process and store the received data, the real-time (or near real-time) streaming data alert systems, the pattern and predictive analytics, etc.

Successful organizations plan for the entire application lifecycle—from development to operation to scaling to end-of-life.

Essentially, this is the ability of the IoT ecosystem, both for any given application and all such applications in general, to grow at the same rate as the predictions—to make them a reality rather than hype. The demand for IoT applications, devices, and services will continue to grow exponentially, and companies with connected devices will need to scale resources accordingly. For example, most IoT platforms let the customers rapidly provision cellular devices for service at volume. Requests are not sent in by humans; rather, automated systems make the provisioning requests, and automated systems process these requests.



The Growth Stall

Many companies run into difficulty after deploying their first few hundred or thousands of IoT devices or, in rare cases, even after tens of thousands of devices. This is not totally surprising because planning for scalability is difficult and involves many factors, both technological and business related.

Sometimes, systems and processes simply reach design capacity, and it is time-consuming and costly to change the architecture of the solution or add capacity. Or the cost of operations becomes higher than expected or planned for, which has a deep impact on smaller companies and startups that are resource constrained. Even seemingly simple tasks, such as generating end-user bills and invoices, can place unexpected burdens on organizations when scaling up for large numbers of devices.

The key issue for businesses caught in this growth stall is that planning for growth was secondary to getting their products and services launched. This is quite common, but it doesn't have to happen. Successful organizations plan for the entire application lifecycle—from development to operation to scaling to end-of-life.

How Big Can IoT Resource Requirements Grow?

The predictions for deployed numbers of devices clearly are enormous numbers. This has created a need to change some of the resources used for IoT applications.

Even before IoT needs became evident, the number of computer systems on the public internet had increased to the point where the internet address numbering method, called IPv4, had been exhausted years ago. The approximately four billion possible IPv4 addresses had been used up, as discussed in earlier chapters.

And, with the ever-increasing number of IP devices, including cellular smartphones that need an IP address, it is no longer possible to consider using stopgap measures, such as the Network Address Translation (NAT), which was introduced to the internet for computing devices.

Thus, IPv6, which was introduced to increase the number of potential IP addresses, is a real requirement for all future deployments. In theory, this range is large enough that it is unlikely to get exhausted for millennia.

Computing resources also can be scalable, particularly if the device traffic and application processing can be stored and processed. New database technologies have been deployed that are far more expandable than the traditional databases used in the past three or four decades for data processing.

Successful organizations plan for the entire application lifecycle—from development to operation to scaling to end-of-life.

Cloud Computing

Cloud computing technologies have provided a scalable solution for storing and processing the data gathered by IoT devices.

Since the numbers of devices are growing rapidly, systems to process the data must grow equally fast. Adding capacity at private data centers is not easy for most companies, since purchasing the physical space, providing for additional power and cooling, increasing the network bandwidth and throughput, installing the computing systems, etc., can take significant effort and time.

Commercial cloud computing suppliers excel at this task. It's their business to provide the compute resources, network bandwidth, and general facilities for exactly this growth purpose. Customers using cloud services can "spin up" resources as needed, in step with their IoT application growth.

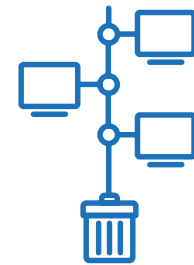
END-OF-LIFE MANAGEMENT

Many IoT devices have an end-of-life that must be managed. The in-service period generally is much longer than the typical period we expect for electronic devices and consumer cellphones today, particularly for industrial applications. But, once the end-of-life of a device is reached, its removal from service must be managed to avoid tying up resources.

For example, in cellular networks, devices have a number that identifies them to the network for operational, accounting, and authentication purposes. In CDMA, this is the Mobile Identification Number (MIN), International Mobile Subscriber Identity (IMSI), or Mobile Directory Number (MDN). In GSM, this may be the IMSI or the Mobile Station ISDN (MSISDN).

These numbers are assigned from an allocated range, or number pool, and create a resource that must be managed. Ideally, these assigned numbers then are re-used when devices are removed from service permanently.

Devices removed from business service still may have a presence on the networks and impact performance if they still are electronically operational. For example, cellular devices used in automotive applications can be removed from service but still could attempt to "register" on the cellular network every time the vehicle is turned on and off.



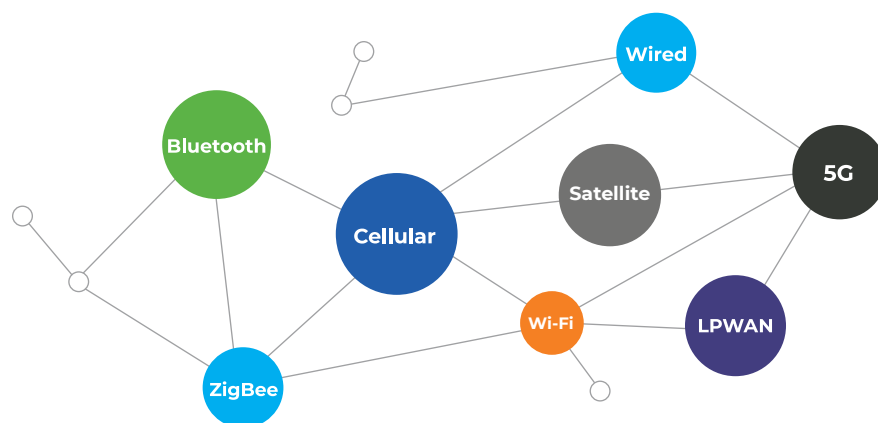
Thus, it is important for devices to have an ability to be turned off—permanently or temporarily—with code in the software of the device that can be executed remotely. This would allow the application servers to properly remove the device from service and, in the case of permanent removal, allow the device resources (such as the numbers in the device) to be re-used for other devices or applications.

SCALABILITY AND CONNECTIVITY

When building scalability into an IoT deployment, selecting the appropriate network connectivity is crucial. The range of available data transport technologies for IoT devices is varied, and new options are becoming available. When planning for scalability, it's important to understand the current choices and what's on the horizon. However, this decision is largely dependent on the type of application.

The first question to be resolved is whether the application is fixed or mobile.

For simplicity, IoT applications can be classified into two categories: those that are fixed (immobile) at one location and those that are in motion while providing the function of the application. These two categories have differing characteristics that affect the specific network selection and implementation for the transport of data from devices.



In fixed location applications:

- The devices are installed at a single location.
- They generally do not move during the normal day-to-day operation of the applications (although they could be re-installed at some other location during their lifetime).
- During this operation, they generally are in a single service boundary.
- The devices often use wired networks in deployments where easy wired solutions are available.
- Wireless networks also are used, however, since network wiring may not be convenient or available.
- The solutions may be hybrid—using short-range wireless to reach a gateway that uses a cellular or wired connection to connect to the servers.

Fixed location devices often are wired. This could be with a Local Area Network (LAN), such as Ethernet using IP protocols. Older deployments used dial-up telephone lines to reach a remote server directly or connect to the internet, and cable modem connections are used where available (also using IP protocols).

In physically mobile applications:

- The devices are installed on moving objects to provide the functionality of the IoT application.
- They physically move from one place to another during the normal operation of the applications.
- During this operation, they often traverse multiple service boundaries (for example, cellular switch boundaries).
- Using some form of long-range wireless network is natural and required.
- In this category, using cellular or satellite networks is quite common.
- For some applications that must transmit while traversing service boundaries, the technology must be a Wide Area Network (WAN) with mobility management.

For short-range data transmissions, where using a wired solution may not be practical, wireless technologies, such as Bluetooth, Wi-Fi, ZigBee, etc., are quite popular. These are common industry standards for which low-cost implementations of the wireless radio and their protocols are available in integrated circuits. The low cost of these short-range wireless technologies enables using them directly within sensors.



These short-range wireless technologies generally are quite limited in range—from a few meters to a few hundred meters. If the data needs to go further, the short-range communication typically is sent to a gateway that then connects to the servers using cellular, cable, or some other IP network transport.

For medium-range wireless transports, typical implementations of IoT solutions use cellular for communication to a nearby tower (generally within a few miles) that then backhauls the data into the internet or a remote server.


When cellular is not available, such as on ocean-bound ships or remote geographies with low human presence, long-range satellite data services provide a global reach for devices to communicate to a distant server program for that IoT application.

Whichever of these two categories the implementation falls into, fixed or mobile, will drive the selection of the network and communications path for the application.

Wired Data Connections

Wired connections typically are used for fixed location IoT applications. For reaching a server, the cost of the transport is “shared” with general internet access. For many device deployments, this is a very low-cost solution, since the ISP generally does not charge a metered rate—i.e., the often fairly low amount of data sent by the devices at a fixed location does not trigger a high transport cost.

With wired connections, the overall service requires an ISP service or another LAN. The quality of the service and general network availability also depends on the ISP. If it is not able to provide continuous service, some mission-critical applications may experience problems with outages.



Cellular and Satellite Connectivity

Service coverage and availability for cellular and satellite generally are excellent. Even in developing countries, cellular service usually is available wherever people live and along major highways.

If cellular is not available, as in truly remote locations such as an ocean-bound ships or in mountain regions, the coverage from satellite data services is excellent, although some of the satellite services may have relatively higher latencies (the time for a data packet to traverse end-to-end) than other technologies. Coverage inside “urban canyons” with tall buildings usually is difficult for satellite data services, but this is where cellular services can excel. If required, a hybrid cellular / satellite device, with multiple radios, can provide truly global data access.

In both cellular and satellite, the cost of the radio can be high relative to the rest of the device, and the radios generally consume substantially more electrical energy to transmit—the communications range is relatively long. For example, it would not be practical to equip low-cost sensors or simple IoT application devices with cellular or satellite transports. These would be far better served by short-range wireless technologies, such as Bluetooth or Wi-Fi.

There is one other concern with cellular technologies—the longevity of deployment is driven by smartphone users. Thus, the technologies evolve relatively rapidly and devices using cellular services must be replaced after some period of time—longer than typical smartphone user turnover, but less than older traditional wired technologies.



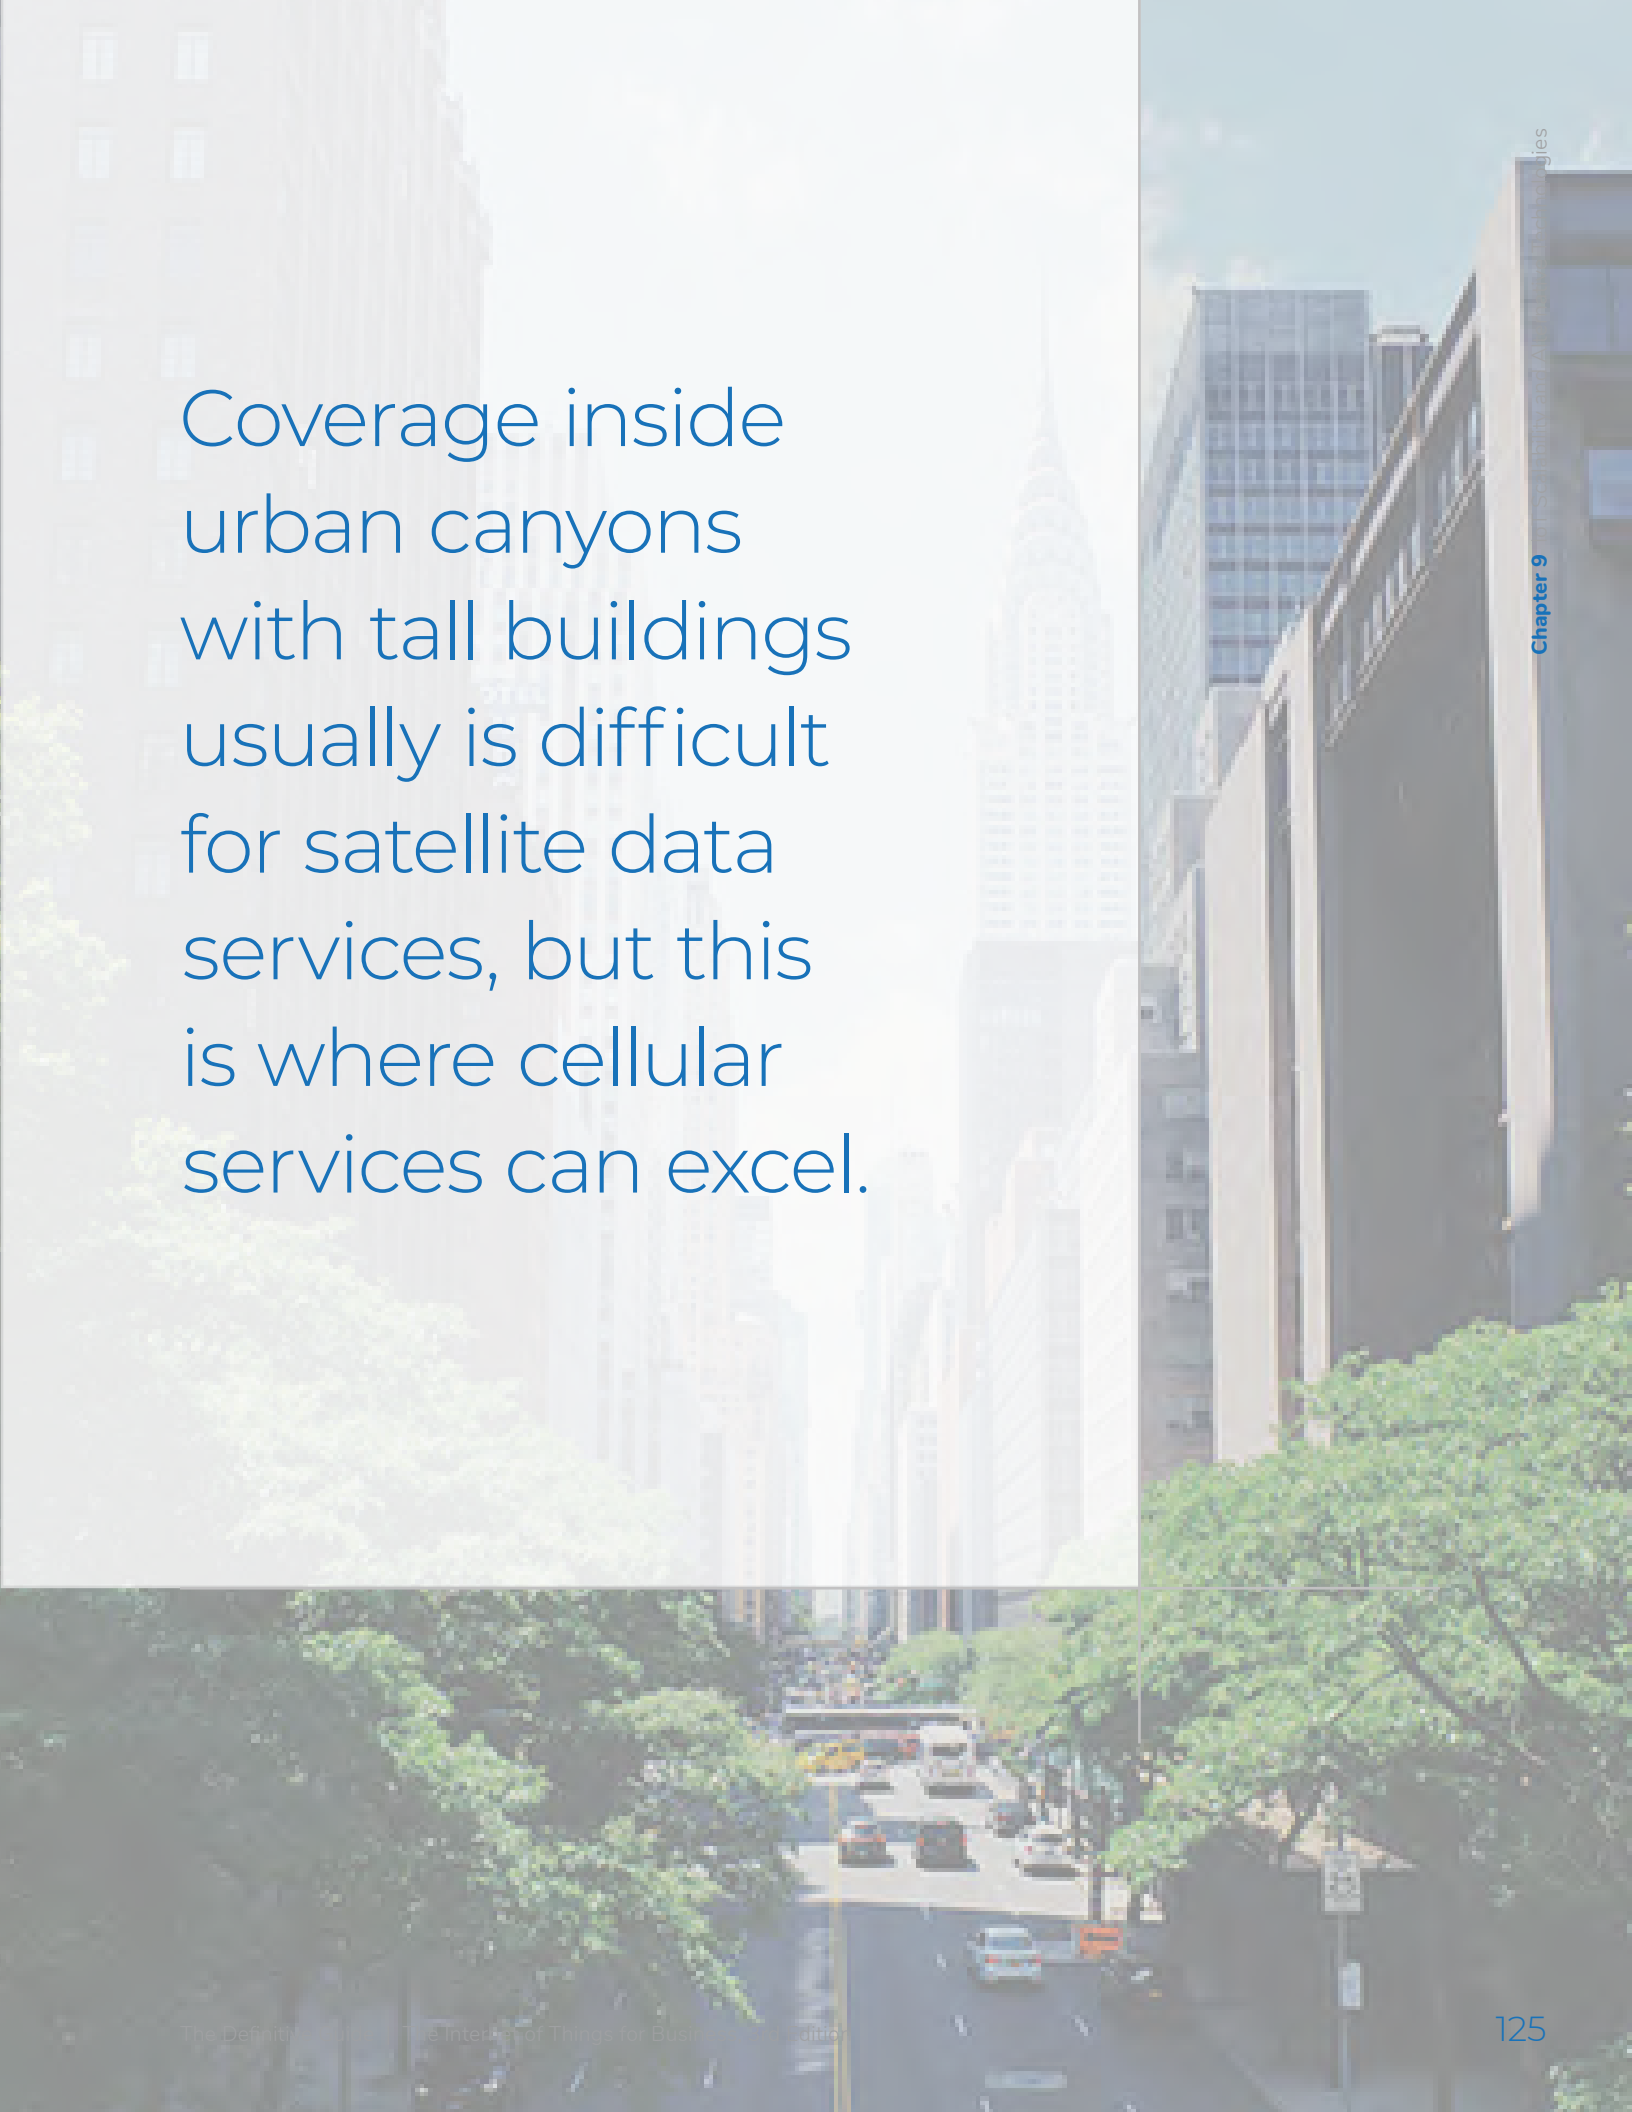
Short-Range Wireless

In many IoT applications, short-range wireless data technologies, such as Bluetooth, Wi-Fi, or ZigBee, are in common use. For certain consumer IoT applications (such as fitness application devices) that only transmit to a nearby smartphone, using Bluetooth and low-power Wi-Fi are common choices. These allow the users to gather data via applications on their smartphone. The need to further transmit the data to central servers for processing is not a paramount requirement but can be done with ease from the smartphone, if necessary.

Short-range wireless is relatively low-energy, so battery-powered devices are designed and deployed easily. In some low-use applications, the battery may last for months or years before it needs to be replaced. This is a key advantage over cellular and satellite applications that require far more frequent energy replacements (for example, using rechargeable batteries that might last a few days).



For many home and business applications, a gateway that provides one or more short-range wireless technologies for deployed sensors and low-power, low-cost, data transmitters are ideal for a number of IoT applications. The gateway communicates to the application servers using cellular or wired ISP connections.



Coverage inside urban canyons with tall buildings usually is difficult for satellite data services, but this is where cellular services can excel.

Low-Power Wide Area Network (LPWAN)

Recently, the need for low-cost, low-power applications that offer longer transmission ranges (between 2 to 20 miles) has seen the development of a number of new technologies and services competing for the large-scale deployment of consumer and industrial IoT devices and applications.

Some of these are the proprietary LPWAN technologies and include the commercial data service networks by Sigfox and its licensees in some countries in Europe and elsewhere (and some cities in the U.S. as of this writing). Similar (but not identical) data transports for IoT include the technologies developed and deployed by Ingenu and Nwave.

The open standards effort by the LoRa Alliance primarily is geared to private network deployments rather than public data networks, although companies also are engaged in deploying LoRa for public access. A number of cellular operators have opted to deploy LPWAN technologies for public access by IoT applications.

The proprietary LPWAN technologies currently use unlicensed wireless spectrum at various standard frequencies. Thus, they may experience congestion and interference, and have technology and data rate limitations that are solved in different ways. For some transports, the data rate and message sizes are low enough that a simple approach to overcome the congestion problems is possible, although the data mostly is one-way (from the device) for low-power use. Others provide more complex data encoding to reach the tower networks, leading to more expensive radio solutions that may work for some IoT solutions, but not necessarily for all.

Finally, the International Telecommunications Union (ITU) has developed, through 3GPP, a set of cellular standards that extend 4G LTE technology for use with low-power, power-efficient, and low-cost IoT radios. The first technology is LTE-M (also called CAT-M), which competes well with the proprietary LPWAN technologies being deployed today.

The 3GPP standards body also ratified the NB-IoT standard for use with IoT applications. Devices and networks using NB-IoT are expected to be deployed in the next few years and will provide alternatives to the services offered by SigFox, Ingenu, Nwave, and LoRa.



A number of cellular operators have opted to deploy LPWAN technologies for public access by IoT applications.

Fifth Generation (5G) Cellular

The first draft of 5G specifications were released in December 2017 but, even before that, work already was underway. The official final 5G standards are not due out till 2020. However, some point releases have been created, particularly for fixed wireless and for the New Radio standards. When finalized, the official standards for 5G services will include a requirement to accommodate large-scale deployments of IoT applications and devices.

Overall, 5G requirements will provide:



The transport of 1000x more data volumes than smartphone users are using today.



More than 10x to 100x the number of connected devices in use today.



Dramatically lower latency (for end-to-end data packets) below a few milliseconds.



Projected 10x longer battery life for low-power devices—up to 10 years.

The 5G specifications incorporate various LPWAN network capabilities for IoT devices. Furthermore, carriers have begun testing in a few markets with a deployment of fixed 5G gateway modems for internet access by home and business owners who now have a wireless alternative to traditional ISP services. These gateways are pre-standards units that may require modification later, once official ITU and 3GPP standards are ratified.

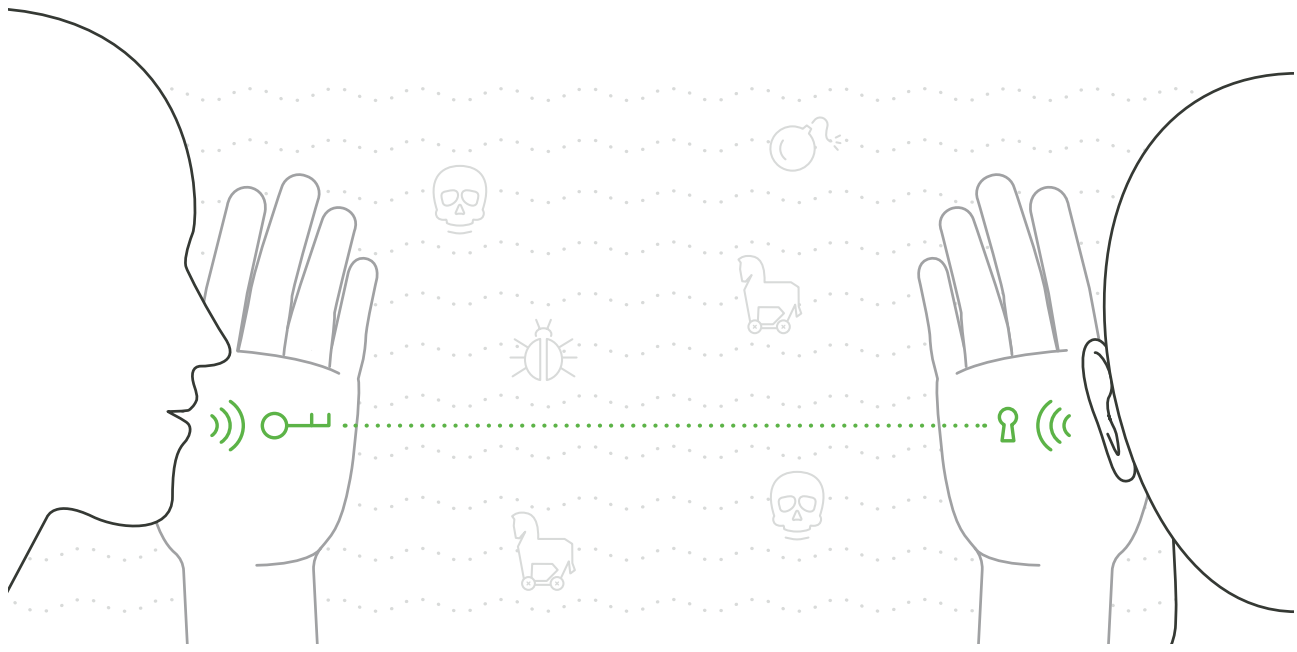
In time, 5G and the IoT cellular LPWAN technologies will be the cellular transports of choice for these low-power applications, along with a base of unlicensed spectrum devices deployed by the proprietary technology providers.



SECURITY, PRIVACY, AND THE INTERNET OF THINGS

130	PRIVACY AND SECURITY
133	INTERNATIONAL DATA TRANSPORT
133	SECURITY OBJECTIVES
135	SECURITY ISSUES FOR IoT
137	RISK MANAGEMENT AND ASSESSING THE IMPACT OF BREACHES
141	ENCRYPTION AS AN IoT TOOL
142	CHOICE OF ENCRYPTION ALGORITHM





SECURITY, PRIVACY, AND THE INTERNET OF THINGS

In her keynote speech at the Consumer Electronics Show in January 2015, former U.S. Federal Trade Commission Chairperson Edith Ramirez noted, “Any device that is connected to the internet is at risk of being hijacked.”

Whether that device is a smartphone, an automobile infotainment system, an automated diabetes monitor, or a GPS-guided farm tractor, specific protections for the security of that IoT device and application must be built into the entire solution.

Traditionally, companies in the financial and consumer markets have been targets for misuse of information stored on their systems—including personal credit information, identity theft, misuse of credit cards by unauthorized persons, personal privacy violations, and loss of corporate intellectual property. The financial losses sustained by these security breaches are in the billions of dollars. While attempts have been made to criminalize such nefarious activities, they continue to occur with increasing frequency and are a serious problem for governments, businesses, and individuals.

Businesses deploying IoT solutions for their customers and themselves are being held responsible for protecting data and devices, as well as corporate proprietary information. Recent media reports of security compromises in the medical and automotive industries have shown that aspects of such IoT deployments can be used for purposes other than the applications for which they originally were designed.

This chapter covers the basic requirements of security implementations and the different methods commonly used to increase the overall security of IoT data and applications.

PRIVACY AND SECURITY

In the context of IoT, privacy is concerned with ensuring that data access is limited to appropriate and authorized parties only.

While using tools, such as data encryption, is an important part of this data securing process, it is just one part of the puzzle. There are other mechanisms and methods to protect privacy (although not just for IoT applications):



Physical access security (for example, secured entrances to data centers).



Security training (for employees on how to secure computers and devices, as well as to understand data safety).

Intrusion detection (for systems that process and store the data) and applying ML / AI techniques to continuously learn the ever-changing intrusion patterns.



Software updates (to implement the latest versions of software for security fixes).



Regular security auditing (helps identify the gaps).

Individuals have an expectation of privacy with regard to their personal data, and it is crucial for businesses to implement relevant security methods. In particular, financial and medical industries have specific governmental regulations that govern their products and services in their respective markets. The new IoT implementations that companies in these industries are deploying may have special testing and certification requirements—particularly in regard to security and privacy issues.

GDPR and Data Privacy

Regulations for Privacy

Many governments are implementing new regulations that require companies to protect the privacy of personal data of individuals in their jurisdictions, with material consequences for mishandling such data. These regulations may require consent from individuals, limitation on downstream use and processing of data, or even specific security standards for handling and storage.

The most notable of these initiatives is the General Data Protection Regulation (GDPR), which was adopted by the European Union (EU) and went into full effect on May 25, 2018. GDPR creates a uniform set of data privacy laws to protect all persons in member states. While it is well beyond the scope of this book to discuss GDPR thoroughly, it should be noted that the penalties for non-compliance can be severe and have the potential to create serious business impacts.

Data will be considered “personal data” if it reveals important information about a natural, identifiable, living person. In many cases, this is obvious. But in the IoT area, other data, such as location data, also may be protected.

Many businesses are choosing to design their products, processes, and systems to meet a single set of privacy requirements and are selecting GDPR as their global standard since it is a stricter set of rules than those in effect in many countries, including the United States.

Companies providing IoT services in the EU must be mindful of how local laws, such as GDPR, apply to them. Services should be designed from the ground up to comply with its requirements, which will necessitate close collaboration between product, engineering, marketing, and legal teams.

Large Territorial Scope

GDPR applies to all companies in the EU that collect or process any personal data of individuals in the EU, regardless of where the data is stored or processed. It also may apply to companies located outside the EU if such companies offer goods and services to persons in the EU (whether or not payment is required) or use personal data of individuals in the EU to monitor their activity or to make automated decisions about them.



Lawful Processing

GDPR states that collection and processing of personal data of EU data subjects is lawful only if it is fair and transparent. Processing is fair only if the collection follows the principles of “Privacy by Design”, the data is protected at each step and kept no longer than required, the individual data subject has the opportunity to give unambiguous and informed consent to the processing, and the data controller only uses the data for the disclosed and agreed purpose. It is transparent only if the data controller, when requesting consent, tells the individual what data is collected, who will have access to or process the data, and what the data will be used for. The consent also must say whether the data will be transferred outside the EU. Companies will need to use clear and plain language in their consent forms, and they cannot condition access to the service on receipt of consent to uses of the data beyond what is absolutely necessary to provide the service. Data subjects also need to be provided with a way to withdraw their consent, to ask that incorrect personal data be corrected, or to ask not to be contacted in the future. Services that allow for public searching of data also must provide a way, in certain circumstances, to “erase” data about an individual upon their request.

Privacy by Design

GDPR obligates companies to follow a “Privacy by Design” approach in building their services and systems. The core concepts are data minimization, meaning collecting and processing the minimum data set necessary for the service, and securing the data against known risks of unauthorized access, deletion, alteration, or loss of availability.

Security and Breach Notification

GDPR requires that companies holding and processing personal data take appropriate measures to protect the data. This may include encryption in transit or at rest, use of firewalls or intrusion monitoring tools, and allowing access only to employees or service providers who have a clear need and who have undergone training. Security also requires the data controller to either securely delete personal data from production systems when no longer needed, or thoroughly “de-identify” it so that it is no longer personal data.

In the event of a breach that could impact the rights and freedoms of EU citizens, the data controller must notify their local authority within 72 hours of first becoming aware of the breach and must notify the affected data subjects without undue delay.

Data is
protected
at each step
and kept no
longer than
required.

INTERNATIONAL DATA TRANSPORT

Along with regulations, such as GDPR, enterprises may need to conform to national regulations in certain countries that prohibit the transport of data beyond their national boundaries for processing and analytics. This data includes the privacy concerns of individuals and companies described above, but also includes other data as well.

Thus, enterprises deploying IoT applications in these countries must understand whether the data must be analyzed in-country, using local servers or commercial cloud facilities, rather than being sent to servers outside national borders.

Today, it is unclear if these requirements consistently include “control plane” messaging that is used to manage the devices (such as cellular registration for roaming units with a home database server) rather than the actual content of the data that they transmit during normal operations. The regulations often are all-inclusive and could lead to operational difficulties if they are violated. It is best to clearly obtain and follow the rules for the country where the IoT devices are deployed. If exemptions are possible (perhaps through anonymization of the data), they should be described, in writing, by the national regulators in order to avoid future issues.

SECURITY OBJECTIVES

There are four basic security objectives that must be met for all IoT security implementations:

- Authenticated sender and receiver
- Sender and receiver accessible
- Trust in the data content
- Confidentiality of information

Authenticated Sender and Receiver

In any data connection, it is important for the sender and receiver of information to be authenticated to each other, regardless of whether the device is the sender (for remote data gathering and transmission) or the receiver (for data and control messages from the server). As a general security principle for transmitting data, the device must ensure that it is sending its information to the correct server, and when receiving data and control messages, it must ensure that the information is coming from the correct server.

Sender and Receiver Accessible

In any network, the sender and receiver always must be accessible when needed. If the network is not functional, or the server is not executing the correct programs and processes to receive the data, the purpose of the application may be lost. Mission-critical applications, such as automatic crash notification or medical alerts, may fail to work properly if the connection is not reliable. The lack of communication itself means a lack of security.

Trust in the Data Content

The accuracy in the content of the transmitted data is essential. If a device does not encode and transmit data correctly, or if the connection is not error-free, the quality and accuracy of the data becomes suspect. Even good data becomes unreliable, and business actions that are taken on the content of the data may not be appropriate.

Mission-critical information is particularly important to keep as error-free as possible. The cost of business actions taken on receipt of incorrect data may be high.

Confidentiality of Information

Finally, the confidentiality of the information must be maintained. Only the correct recipient should have access to the transmitted data because it may contain proprietary or confidential information. Indeed, privacy laws in many countries require extra care with information regarding individual citizens. For example, in the U.S., the Health Insurance Portability and Accountability Act (HIPAA) provides specific rules for individually identifiable medical information. And, of course, the GDPR requirements mentioned earlier are extremely important for applications that operate in the EU or interact with its citizens.



SECURITY ISSUES FOR IoT

Security risks generally can be recognized and understood, and the implementation of security methods should be incorporated in the IoT device and software associated with that application during the design phase. The concept of “Security by Design” is essential. The nature of these new deployments brings new complexities to creating secure solutions.

Most obvious holes in security can be resolved quickly and efficiently. In general, the potential for problems can be managed with confidence with the chosen security methods. However, it is vital to recognize that risks can never be completely eliminated, and there is no single security solution for all possible security requirements for all applications.

Thus, it is critical to assess the level of security implementations that are appropriate for different kinds of data. It is imperative that this assessment be done early—during the design of the application and devices, not as an afterthought once many devices have been deployed.

Before choosing how to secure the application, there are a number of issues to be considered:

- Authenticating presence on multiple data transport networks.
- Authorization for multiple types of services.
- Scaling to manage the large number of devices in IoT solutions.
- Automation for application functionality.
- Long lifecycles for deployed devices and applications.
- Implementing security updates in remote devices.

Multiple Networks

Some IoT devices operate in more than one transport network or technology for redundancy or hybrid solutions. In these devices and solutions, security may be more of a concern in one particular network compared to the others. For example, a short-range wireless technology, such as Wi-Fi, can have quite a different security threat vector and potential for breaches compared to a long-range cellular service.

Multiple Types of Services

Applications and devices may use multiple services, where the required authorizations for allowing a device access to a particular service may differ from one application to another. The authentication mechanisms also may differ, and developers must minimize the risk of a less secure service authentication system from allowing a device to be compromised.

Scaling Growth

In IoT deployments, there are predictions of explosive growth in the near future—billions of potential devices within the next 5 to 10 years. Thus, in any application where a security issue exists, the overall problem could be magnified greatly by the large numbers of devices that may be affected. This could result in network and data security issues that are difficult to solve, since replacing all the compromised devices could be extremely difficult, perhaps impossible.

Automated Functionality

In many IoT applications, the data is acted upon by automated programs that process the received data and take business actions based on the content. If the transmitted data is compromised, any simplistic responses or automated functions to that compromised data could cascade into greater difficulty. If some set of devices transmit excessively due to a program error, the servers processing that incoming data could overload and not provide a timely response to the device transmissions. Simplistic retry algorithms in the devices may create a data storm as a result.

Long Lifecycles

Unlike handsets used by people who change them every few years, IoT devices—particularly in industrial applications—may be deployed for many years and operate continuously over that time. Often, the devices use electrical power rather than batteries (unlike handsets that shut down when battery energy is depleted), and the IoT devices could continue to use the networks for years. Devices with compromised security could stay operational for lengthy periods.

Remote Updates

Therefore, it is essential to plan and design for device updates via over-the-air (OTA) notifications—not just for application feature updates, but also to update the security implementations within the devices. When a device security breach is sufficiently critical that the device firmware must be updated, the ability to reprogram the functionality remotely is vital. The devices may be in inaccessible locations or a large number of devices must be modified rapidly. However, the ability to perform OTAs also introduces additional risks.

RISK MANAGEMENT AND ASSESSING IMPACT OF BREACHES

For some data, the issue of security may not be as critical. For example, if an IoT device is collecting temperature information from a residence for monitoring (not controlling) purposes, the security needs for this data to be protected is not as stringent as a device that collects and transmits credit card information for financial transactions.

Thus, the effort and level of security implementations, as well as the methods, necessarily differ for these two examples. One may require anonymizing the data source for simplicity and privacy, and the other may require strong data encryption to prevent unauthorized access to the data.

It is important to remember that even if we could determine all possible threat vectors, the cost of designing preventative measures to counter every threat might be prohibitively expensive. In all IoT deployments, it is important to assess the potential for damage caused by a security breach and implement security solutions accordingly.

To start, ask the following questions:

- If a single device is compromised, can it be used to compromise other devices? The data transport used by that application? The remote application servers? That entire application?
- If an application is compromised and misused, what impact does that security event have? Is it life threatening to one individual? To more than one individual? An entire population in a region?
- Can a data content breach cause financial harm to an individual? More than one individual? The entire set of people depending on a particular IoT application to function well?
- How quickly can the specific breach or intrusion be detected? Is it using a well-known target mechanism (such as might exist in a widely used cellular device operating system)?
- Can a compromised device, or set of devices, be isolated from the application rapidly?
- If the data is personal information, does the breach violate the GDPR (or similar regulation) and necessitate the need for timely and proper compliant handling of that breach?

The background of the page is a photograph of an industrial facility. It features various pieces of machinery, including large cylindrical tanks and pipes. Several bright blue cables are visible, running across the scene and connecting different parts of the equipment. The lighting is bright, suggesting an indoor or well-lit outdoor environment. The overall tone is professional and technical.

In all IoT deployments, it is important to assess the potential for damage caused by a security breach and implement security solutions accordingly.

The opportunities for implementing security best practices occur at different points, with differing capabilities, in the IoT data chain. Developers should assess the opportunity for implementing security best practices (authentication, encryption, breach detection, etc.) at every point of that chain during the design of the application—adhering to Security by Design principles.

For example, the source of data could be a sensor. These are not likely to be compromised easily, since they are very specific to their function, but they still need protection. However, because of the simplicity of such sensors, it often is difficult to implement a security solution for them.

Regardless, a compromised sensor could be used to inject false data into the application, where an incorrect action might be taken by a server or human at the remote end of the chain. Data analytics systems are useful to appropriately deal with false data.

A more complex source device, such as a multi-technology gateway connecting to multiple types of sensors, or a cellular modem, offers more opportunity—both for breaches to occur, as well as a location in the chain for implementing a good security solution. For example, a gateway device could have the compute capacity to implement encryption algorithms, thereby securing the content further along the chain to the servers that receive the data.



In general, the “closer to the device” that security best practices can be implemented, the less impact a security breach can have on the overall application. Indeed, it could be possible to isolate a subset of devices (or applications) that are breached if the impact might be significant to the entire application (or network) as a whole.

Each business and its IoT application implementations will require its own risk assessment to determine the relevant security needs. And organizations have to understand the trade-offs they make up-front. It is simply impossible to determine all possible methods by which all such IoT data applications could be compromised.

Even if we could determine all possible threat vectors for a particular application, the cost of designing detection and preventative measures to counter every threat might be prohibitively expensive for that application. The best we can do is understand and minimize the risk as best as we can up-front, then design the devices and application processes to be as easily updatable as possible.

While server programs and accessible elements of the data chain can be updated more easily, the ability to re-program devices using OTA updates is key to ensuring that the impacts of security breaches can be contained and repaired. All device developers should look to implement OTA updates if possible and practical, even if it potentially increases the cost of the device—for example, adding on-board memory to support multiple “images” of firmware for updates.



ENCRYPTION AS AN IoT TOOL

One of the most basic technological tools to secure the content and data in an IoT deployment is to encode the data so that only the authorized recipient (whether it is a program or human) can decode the data.

After the data is gathered and transmitted by the remote device (or is sent by the server to the device), the content can be encrypted at various points along the network, as well as when the data is stored.

The basic goals of encryption are to provide:

- **Proof that the sender is valid**—Encryption can make the data's source relatively irrefutable. Techniques such as electronic signatures on a document can be a sign of irrefutability. Proof of who sent data is crucial so that a hacker doesn't steal a session and then pretend to be that user, which is called spoofing.
- **Proof that data was not altered**—Encryption functions can be used to ensure that a change to the data renders the content unusable to an unauthorized recipient.
- **Proof that data cannot be read by a third party**—Encryption protects data from being read in transit or upon receipt, except by someone (or a process) with the correct decryption method.

Data encryption can protect the content in each of these areas to different levels, depending on the need and the specific type of encryption that is used.

Weaknesses in Encryption

No encryption method is perfect. Depending on the computing power available at a particular location, or the processing time used by the encryption method, the algorithm may be weak or strong. In some low-cost sensors, it may be impossible to implement encryption if the processor within the sensor simply cannot perform the task—the developers must look for opportunities further up the data transmission chain to implement the encryption.

Strong encryption may seem impossible to break but applying enough compute resources to the task could reveal weaknesses that allow the data to be decrypted by unauthorized systems or people. Additionally, future advances in quantum computing will make breaking the encryption even easier.

Indeed, bugs may be discovered in the method itself, or in the particular software implementation. A recent example is the Heartbleed security bug in OpenSSL discovered in 2014. This affected about 17% of the world's web servers and, potentially, allowed encrypted data to be read. Patches were made to OpenSSL, and a majority of web servers have since been updated.

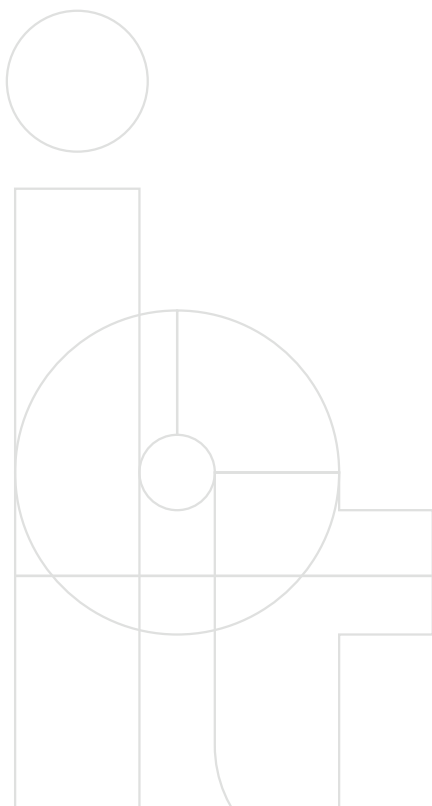


CHOICE OF ENCRYPTION ALGORITHM

It is beyond the scope of this book to describe or recommend a particular encryption algorithm. The specific requirements of the IoT application or the computing power available at a place in the data chain may drive a preference for a particular algorithm.

Security experts can provide guidance for selecting an approach and should be consulted during the design of the application.

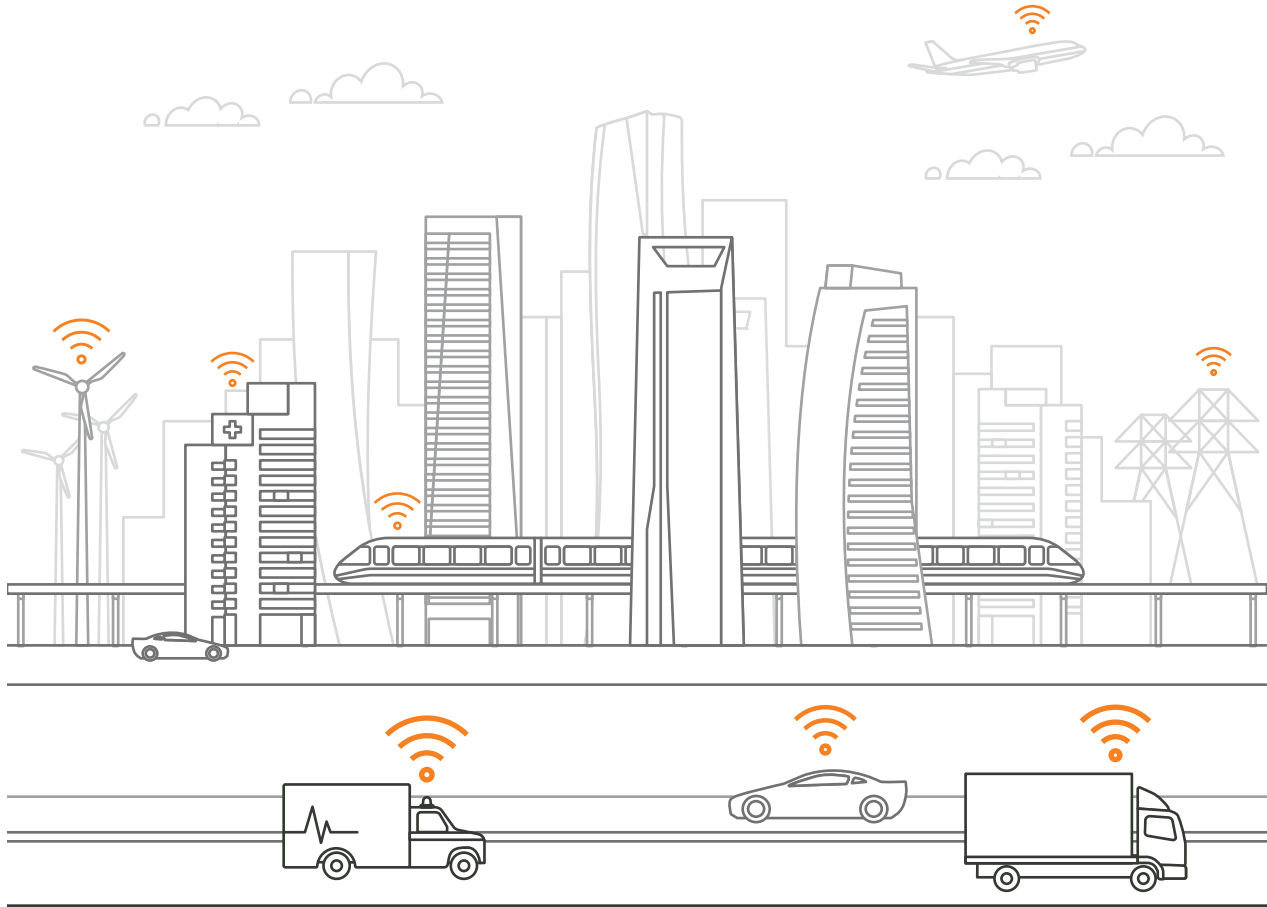
The Information Technology (IT) departments at each company may have specific encryption and security requirements, as in the use of Virtual Private Networks (VPN) to transport data into its servers for processing and storage.



IoT USE CASES

145	RENEWABLE SOLAR ENERGY
146	AUTOMOTIVE
149	HEALTHCARE
152	SMART CITIES
154	FINANCIAL / INSURANCE





IoT USE CASES

We are a connected world. We now communicate with machines, with systems, with people. That instantaneous connection, driven by real-time data, opens up opportunities for businesses and individuals in ways not even imagined just a few years ago. With IoT, the opportunities are endless. Below, we outline just some of the use cases in which IoT can influence outcomes, enhance business efficiencies and opportunities, and improve lives.

RENEWABLE SOLAR ENERGY

Companies operating in the most remote locations, with products purpose-built for off-grid, rural, and often hostile environments, require a reliable global mobile network that provides consistent connectivity worldwide to enable effortless remote monitoring of solar energy systems. To overcome these many life-critical, energy-delivery issues, there is significant need for reliable GSM and CDMA connectivity to deliver functional, energy saving solutions.

The World Energy Outlook estimated that 1.2 billion people, equivalent to 16% of the global population, do not have current access to electricity, with many more people living with an electricity supply described as poor quality or unreliable. More than 95% of those living without electricity are in Sub-Saharan Africa and developing Asia, 80% of which live in rural areas.

For example, BBOXX, a UK-based solar energy provider, using IoT technologies, has developed solutions to provide affordable, clean energy to off-grid communities in the developing world.

To address these issues, the solar-powered BBOXX system is deployed on a simple plug-and-play basis, without the need to reconfigure to use local network settings. By working with a carrier-agnostic and technology-agnostic partner, BBOXX installed a global Subscriber Identity Module (SIM) at the point of manufacture, reducing both supply chain costs and deployment time.

Advanced IoT connectivity enables energy provider networks to overcome critical issues, such as interrupted services, security breaches, and high implementation and maintenance costs.

Within the residential and commercial solar sectors, SolarEdge, an Israeli company with headquarters in the United States, provides solar solutions for homes and businesses. Company products include power optimizers, solar inverters (DC to AC inversion), and cloud-based monitoring solutions.

With a presence in 13 countries, all with somewhat different connectivity options and providers, secure and reliable connectivity was integral to the company's success. SolarEdge chose a carrier-agnostic and technology-agnostic IoT partner for deploying cellular connectivity, as well as for the corresponding management solution. In dealing with a seasoned, professional connectivity solution provider, SolarEdge was able to secure instant global connectivity and management oversight; advanced revenue-grade metering; and robust, real-time troubleshooting. Visibility into devices improved, operational efficiency rose, customer service radically upgraded, and inventory and loss management became transparent and easy to administer.

AUTOMOTIVE

Fleet

Fleet managers need real-time intelligence to solve transportation issues before they become costly mistakes. They require a connected data transport solution, combined with IoT analytics, to reduce time-to-market processes, resolve troubleshooting issues, and bring down the total cost of ownership. It is only with real-time business intelligence data that retailers and manufacturers can acquire a comprehensive view of their transportation ecosystem.

The global trucking industry is undergoing enormous change. Older vehicles are being replaced with “smart trucks” using IoT systems with cellular and satellite communications technologies to transmit essential information for management of fleet operations.

Commercial fleet programs are becoming more complex—beyond the original simple needs to manage inventory, location, routing, and fuel costs, they now face requirements for mission-critical reliability, cross-region connectivity, and innovative market differentiation.

Long-haul fleet management providers require onboard computing and fleet communications to deliver better business outcomes. On-board solutions require highly reliable, real-time, always-on cellular network connectivity, which might require multiple carriers to meet the full-coverage needs of fleet customers wherever they reside and drive.

The market for global fleet management solutions continues to expand quickly. Data from Transparency Market Research estimates that the fleet management solutions sector will rise in valuation from US\$12.5 billion in 2015 to an expected US\$92 billion by 2025, with a CAGR of 22.6% for



The fleet management solutions sector will rise in valuation from US\$12.5 billion in 2015 to an expected US\$92 billion by 2025.

the period between 2017 and 2025. In North and Latin America, according to Berg Insight, connected fleet management will grow from 5.8 million rolling units in operation to more than 12.7 million by 2020. Additionally, Berg Insights see similar European growth in the sector, expanding from 5.3 million units in 2015 to an estimated 10.6 million by 2020.

As the commercial fleet sector becomes even more competitive, fleet owners and operators are seeking more reliable connections—with flexible rate plans and seamless coverage—across many geographic areas and remote locations. In Europe, seamless cross-border operation for the trucking industry is an essential requirement, as trucks continually traverse national boundaries.

Connected Car

Car makers have been honing in on the monetization of vehicle data. This desire to monetize involves data from internal applications (diagnostics, customer relationship management, marketing) and external applications (usage-based insurance, traffic information). This requires more complex systems for extracting, storing, normalizing, and preserving or deleting data.

Recent surveys estimate that 75% of all new cars shipped globally by 2024 will be equipped with wireless connections. Twenty years after General Motors shipped its first connected Cadillacs, car companies have learned that it is not enough to build a telecommunications module into the vehicle. The entire system supporting that connection has to be properly designed and maintained for the connected car proposition to be viable.

Today, connected cars are using IoT to connect everything from engine diagnostics, to GNSS location data, to actual driving behaviors, to infotainment systems. Modern connected car systems are expected to be always-on and, increasingly, will be asked to support autonomous driving and safety applications, such as collision avoidance. A reactive system, like OnStar's original offering, is no longer sufficient. Today, 80% of connected cars are using technology that is more than a decade old and hardly suited to automakers' rapidly changing needs, let alone those of the customer.

Wireless connectivity, today and into the future, will be expected to maintain continuous connections with increasing demands on the communication of vehicle data and software updates; secure both real-time and historical vehicle data; provide customizable end-user dashboards so as to share data from the vehicle tracking system with end users and customers; and greater insights leading to a safer and more efficient driving experience.



The future of fully autonomous driving is interwoven with the communication technologies in development to make this a practical reality. While the requirements of instantaneous action, such as accident avoidance, means that the cars must process sensor data extremely rapidly—local to the vehicle—other capabilities will be enabled by faster cellular technologies, such as 5G.

For example, updates for general traffic conditions beyond the range of vehicle-to-vehicle (V2V) radio technology, as well as dynamic updates for road changes (repair work, hazards), can be enabled by the advent of faster cellular technologies.

Leasing / Ride Sharing / Asset Management

In many parts of the world, transportation is a huge hurdle. Traffic, costs, and vehicle availability all come into play. To alleviate some of these issues, ride-sharing companies are rapidly expanding globally, trying to fill the need of moving people in urban areas. In fact, revenue from the global ride-sharing sector is approaching US\$57 billion and is expected to show an annual growth rate of 16.5%, resulting in a market volume of US \$106 billion by 2022.

In many countries, drivers want to lease vehicles so they can create their own ride-sharing business driving for a specific brand. This creates a significant risk of loss of the leased asset due to theft, lack of payment, or hijacking. Additional loss can come from misuse of the vehicle.

In order to protect their investments, ride sharing companies need to track all their vehicles, many of which are leased to individual drivers. Companies need data on driver performance and access to vehicle metrics regarding whether the leased vehicle is being used for another service or whether the driver is paying leasing fees. In such cases, ride-sharing companies also need the ability to remotely disable the vehicle before harm can be done to a company or its reputation.

Key to all this is reliable monitoring and tracking connectivity so that data is collected in a timely manner. Poor quality of IoT devices and slow response times to problem management scenarios are situations that also need attention.

Today, ride-sharing companies can install a tracking device in their vehicles, which ensures constant monitoring of vehicle location; insights to driver performance metrics; vehicle metrics; auto-immobilize functionality if the driver is late paying leasing charges or if the vehicle is reported stolen or tampered with; or to halt activities if the driver is using the vehicle for unauthorized purposes. This comprehensive IoT asset management solution allows ride-sharing enterprises to retrieve any vehicle operating outside of company guidelines and to protect its investment in a costly asset.



HEALTHCARE

The healthcare industry shows great promise as IoT-driven systems and applications are improving access to care, increasing the quality of care, while, at the same time, reducing its overall cost. Today, it is one of the fastest growing IoT sectors, with a large number of startups developing new medical sensors, transporting the data to care providers, and achieving the desired health improvement outcomes.

It has been estimated that 40% of the global economic impact of the IoT revolution will occur in healthcare, more than any other sector. And IoT-driven companies can gain a competitive edge in that sector—specifically in areas such as user experience, operational costs and efficiencies, and global expansion.

Cellular connectivity and IoT solutions enable medical device manufacturers and healthcare providers to achieve the highest levels of patient engagement and medical adherence, with the lowest TCO, regardless of global location.

Here we look at several use cases within the healthcare market in order to show the progress, as well as the possibilities that exist going forward.

Patient Monitoring

Estimates show that more than 200 million people in the EU and the U.S. suffer from one or several diseases that may benefit from some type of home monitoring.

New IoT technologies are changing the way health services are delivered, allowing recipients to remain in their homes to receive care and avoid costly hospital stays. Companies are expanding in-home services, providing solutions for independent living specifically tailored to serve aging and disabled populations nationwide. With that in mind, more and more state and federal healthcare agencies encourage in-home care programs as a vital way to deliver efficiencies and reduce costs.



SimplyHome designs and installs wireless technology products and related healthcare-focused services. The company deploys a cost-effective IoT cellular solution that provides connectivity for its services, regardless of the patient's location. Its systems proactively alert patients and caregivers to changes in behavioral patterns by communicating with multiple sensors to observe activities of daily living.

Text, email, or phone alerts can be generated by a single event, an intersection of multiple events, or by inactivity. Components, such as motion-sensors, door / window contacts, and bed pressure pads, alert caretakers to falls, wandering, or changes in sleep patterns. The IoT-enabled SimplyHome system helps residents remain independent with environmental controls that operate beds, lights, TVs, doors, and more via tablet or voice activation.

Medical Adherence

According to research by the World Health Organization (WHO), the benefits of medications used to fight disease are not fully realized because close to 50% of patients do not adhere to medicinal intake guidelines. Reasons for not taking medicines on a regularly needed basis are plentiful, running from lack of funds to sub-optimal healthcare literacy to communication / language barriers to just plain forgetfulness.



As another example of IoT serving the healthcare sector, Wisepill is a leading provider of medical adherence management solutions and the creator of the Wisepill dispenser, a pillbox that uses cellular and IoT technologies to provide real-time medical management solutions. The pillbox, designed to work in diverse environments, has a rechargeable, longer-life battery, which allows the device to be used for extended periods without the need of an external power source.

Patients in developing countries or in hard-to-reach rural areas cannot travel easily to far-off clinics. Additionally, many places have a severe shortage of medical professionals. With IoT connectivity, Wisepill enables clients, pharmaceutical businesses, doctors, and healthcare organizations around the world to improve medication adherence management. The combination of an experienced IoT solution provider and Wisepill provides patients with the peace of mind from knowing that if they miss taking their medications, there would be a reminder to maintain their medicinal intake schedule.

By continuing to apply new cost-saving IoT technologies, and leveraging economies of scale, Wisepill is providing affordable adherence solutions, assisting millions of people, regardless of where they live.

Blood Banks

Blood units represent a critical aspect of healthcare. Yet, blood units often get wasted due to the inability to store them under appropriate conditions.

The principal goal of an IoT technology-driven blood bank management program is to optimize the effectiveness of a blood bank. A successful program involves increasing awareness about best practices; reducing the likelihood of blood samples becoming unusable; minimizing blood loss; improving blood availability; continuously educating clinicians; and standardizing operations through workflows.

For example, in India, the business case for an IoT-enabled blood bank monitoring solution rests on the following goals:

- Monitoring blood bank refrigerators on a 24x7 basis and storing relevant data.
- Alerts of temperature variance outside a set range.
- Use of transparent monitoring network (single pane of glass).
- Reduction of paperwork.

The IoT-based blood bank improvement program includes both a measurement of how well the program meets its goals and also demonstrates a commitment to data driven reporting. This, plus additional functionality, provides the insights needed to initiate and preserve blood bank management that will save many lives.

Wisepill enables clients, pharmaceutical businesses, doctors, and healthcare organizations around the world to improve medication adherence management.



Device Monitoring: Defibrillators / Heart Monitors / Pacemakers

Doctors and hospitals need a secure way to establish connectivity and transmit data from automated external defibrillators to a cloud application. Today, IoT-connected medical devices can monitor and analyze data coming from the patient in real time. And should an event occur, IoT-enabled defibrillators, for example, provide verbal and on-screen instructions in delivering chest compressions. Some advanced defibrillators even can deliver an electrical shock to a patient's heart. Heart monitors send alerts to both wearer and doctor in case of irregularities. In all these cases, device connectivity, with real-time data, literally, is a life and death issue.

With an IoT-enabled solution for healthcare devices, real-time data, along with alerts and reports, can save lives.



SMART CITIES

Virtually every aspect of city operations can be made smarter through IoT—from embedded roadway devices to advanced lighting to waste management. That means there is unlimited potential for IoT providers to deliver a variety of solutions to meet the ever-increasing demands for efficiency and cost reduction.

The human migration to cities now is a global trend that research indicates will continue for the foreseeable future. While this shift has enhanced the economic well-being of millions, it also has placed incredible demands on infrastructure and threatens the quality of life of the inhabitants of large, ever-growing cities.

Compounding these problems is the fact that tax bases and budgets do not match the ever-growing needs. This is where IoT solutions can have a significant and immediate impact.

A “smart city” may sound futuristic, but at its heart, the idea is quite simple and traditional—smart cities bring together current and new technologies, infrastructure, and government to benefit people's quality of life.



Smart cities
bring together
current and new
technologies,
infrastructure, and
government to
benefit people's
quality of life.

Smart city solutions introduce tremendous new capabilities, giving municipalities the ability to remotely monitor, manage, and control devices using IoT technology. These tools help citizens create new insights and actionable information from massive streams of real-time data.

IoT enables traditional cities to become 'smart' by incorporating ecosystems that offer remarkable efficiencies, cost savings, and advanced resource management via automation and connectivity.

FINANCIAL / INSURANCE

The old school vision for asset insurance still is prevalent, but the IoT is changing the landscape in a most disruptive way—all for the good.

In the recent past, for example, vehicle telematics that insurance companies cared about involved risk assessment, vehicle performance, reports, mobile apps, and APIs. Today's insurance provides a multitude of new views into driver and driving characteristics. With sensors and devices absorbing data at an unprecedented rate, insurance now also covers accident reconstruction, false claims identification, overall claims management, driver coaching, alerts and notifications, actuarial support, vehicle immobilization, asset protection, usage-based insurance (UBI), and a lot more.

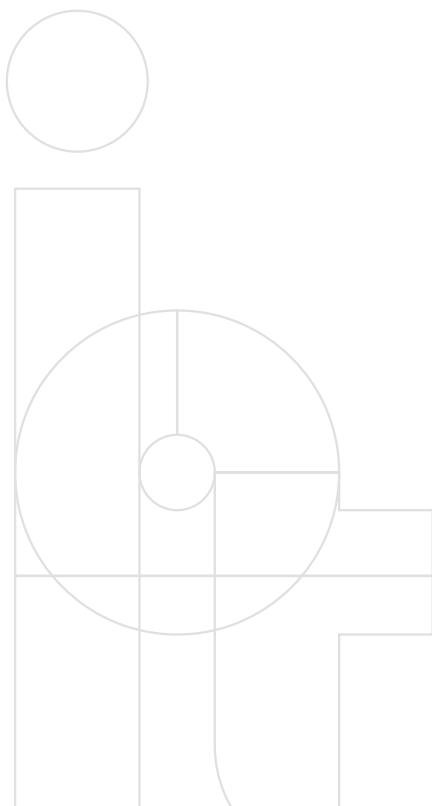


As our world becomes increasingly digital and connected, customer and business expectations for insurance are evolving. A platform-based, highly scalable solution enables insurers to offer value-added services that simplify the insurance process and radically improve customer satisfaction.

- **Claims Management:** Quickly process claims and identify possible fraud. Reduce fraud and claims while speeding up the entire claims process.
- **Customer Management:** Maximize customer value through targeted up-sell, cross-sell opportunities. Attract more low-risk customers.
- **Renewals Management:** Identify customers with high propensity to lapse for targeted collection. Increase customer retention.
- **Sales Force Management:** Identify agents with high potential.
- **Pricing & Risk Management:** Conduct risk-based pricing for better profitability. Gain a higher percentage of low-risk drivers. Reduce underwriting costs. Provide customer premium savings.

The IoT enhances the interaction frequency with customers and provides value through information and knowledge creation.

Bottom line—IoT data usage can impact services by providing insights to risk assessment, loss control, driver behavior, product pricing, and much more.

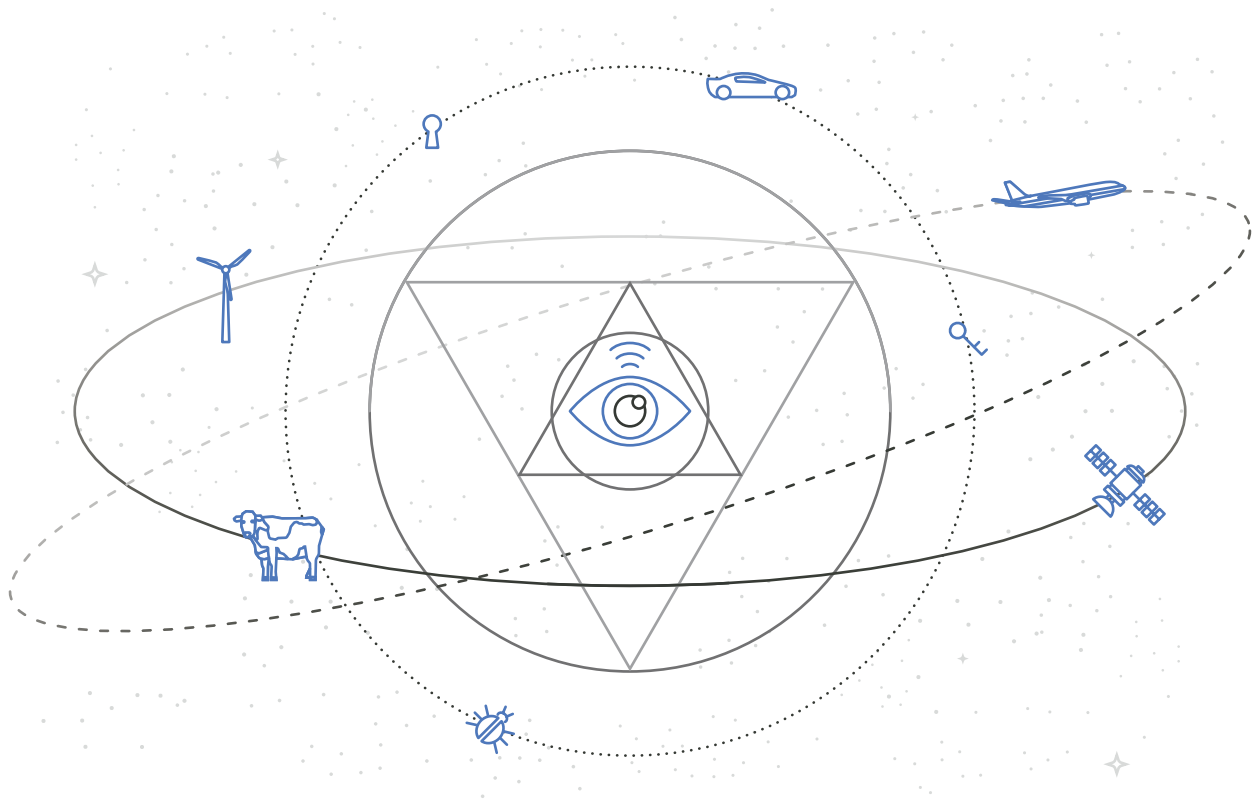


THE FUTURE OF THE INTERNET OF THINGS

- 158 IoT WILL POSITIVELY AFFECT ALL MARKETS
- 158 IoT WILL COME FIRST
- 159 HOMES WILL GET SMARTER—AND MORE CONNECTED
- 160 ENTERPRISES WILL SPEND MORE
- 160 STANDARDS WILL IMPROVE
- 163 SECURITY CONCERNS WILL CONTINUE
- 164 OVER-THE-AIR (OTA) UPDATES WILL BECOME THE NORM
- 165 PRIVACY CONCERNS AND GOVERNMENT REGULATIONS
- 166 IoT VALUE REALIZED THROUGH DATA ANALYTICS
- 166 THE FUTURE IS NOW



**GOOD
NEWS
IS COMING**



THE FUTURE OF THE INTERNET OF THINGS

Imagine a future where enterprise customers and consumers will ask product companies about the sensor capabilities accessible via a mobile app when purchasing a new appliance or car. In the future, not only will cars talk to each other, but people will wear clothes connected to the Internet, reading glasses will be connected to provide additional context to enrich the user's experience, and more than one-half of the Internet traffic to homes will go to appliances and devices and not to children's video games.

Sound unlikely? We think this future may not be too far away. Given the tremendous growth and change that is taking place in the IoT industry, the future certainly is hard to predict, but here is what we see taking shape in the near future.

IoT WILL POSITIVELY AFFECT ALL MARKETS

IoT applications and solutions already are positively impacting many industries. This trend will continue in all markets as businesses and consumers find new opportunities where the ability to measure remote products, gather data, and analyze the transmitted information can improve the overall success of the product deployment.

The return on investment for IoT solutions will be a combination of direct operational cost savings of service revenue for products, better interaction with end-users' needs and requirements, visibility into operational product issues, and eliminating the need for product recalls.



IoT WILL COME FIRST

The most competitive companies, products, and solutions will be those built around the concept of “IoT First”. This means products will be designed from the outset with access to connectivity and data via the IoT as a primary consideration, and enterprises will begin planning projects and building systems with IoT foremost in mind.

Today's typical approach of retro-fitting connectivity to an existing product or service will continue to occur, but those initiatives will not generate as much value as IoT-First projects.

HOMES WILL GET SMARTER, AND MORE CONNECTED

As the cost of sensors, processors, and networking goes down with volume, and consumers are increasingly aware of the benefits of the smart home, IoT home automation will provide greater peace of mind through security implementations.

Remote access is the greatest selling point for IoT technologies these days. Consumers can check the status of their home from their smartphone—for example, consumers can ensure that the front door is locked or remotely view security cameras.

Security systems can push alerts directly to consumers, as well as security agencies, in the case of a break-in or abnormal activity inside the home. Many people see the value of such systems if they often are away from home. A quick check can provide a sense of control even if they are thousands of miles away.

Another important aspect of the connected home of the future is in the area of energy usage. Consumers can automate lighting, temperature, and home irrigation systems remotely. Home automation devices and sensors identify rooms that are occupied and adjust HVAC systems and lighting to ensure energy conservation. Self-monitoring appliances with IoT technologies can determine changes in operation due to potential maintenance issues. Sensors in outdoor landscaping measure the current moisture saturation level of the ground and adjust the sprinklers. Smart meters provide utilities with information about energy usage patterns to assist in planning power consumption across a designated area. Discounts and incentives are given to consumers to change usage patterns to benefit a community. The end result of the connected home is increased safety, convenience, and freedom from mundane decisions that allow consumers to enjoy more of their life.



ENTERPRISES WILL SPEND MORE ON IoT

Currently, many think of IoT primarily in terms of consumer devices and applications, but industry growth shows enterprises will be spending far more on IoT than consumers.

McKinsey forecasts that of the \$12 trillion in economic value generated by IoT in 2025, the majority (70%) of this value generated from IoT will be from B2B deployments.¹ Furthermore, IDC expects that by 2021, more than 70% of the top 2,000 global corporations will be investing in connectivity management solutions. Within the same time frame, IDC also expects that 75% of enterprises with a positive IoT ROI will use tactical analytics applications to reduce operating costs, while about 25% of companies that invest in decision architecture will increase their revenue share.² Just about any way you view it, IoT can make a significant impact on the bottom line.



IoT STANDARDS WILL IMPROVE

The IoT space has undergone rapid growth in a relatively short amount of time, so it's no surprise that standards for IoT still are in active development and not yet widely adopted. With the proliferation of devices, sensors, connectivity technologies, and IoT platforms that need to be integrated to organizational systems, it is challenging to find technology- and vendor-agnostic solutions that work well with each other.

However, interoperability is key to ensuring the long-term success of IoT initiatives, so companies deploying solutions need to be careful when going down the path of a vendor-specific implementation that may not play well with other suppliers. A lack of standards can lead to issues, including legal implications for the accuracy of data, safe harbor issues between nations, or liability in accidents with autonomous vehicles.

¹ The CD "The Internet of Things: Mapping the Value Beyond the Hype," McKinsey Global Institute, June 2015.

² "IDC Futurescape, Worldwide Internet of Things," IDC, 2018.

Vendors, solutions providers, and distributors also will undergo changes as a result of deficiencies in current standards. Beyond the consolidation in the industry that is typical of any fast-growing, dynamic sector, IoT companies will move to providing fully enabled solutions instead of just products and basic services. Companies will see a shift from one-time hardware sales to “Product-as-a-Service” sales, where end users rent (or hire) the capabilities of the product without incurring a large initial expense. This provides an opportunity for enterprises to receive stable, long-term revenues, increase the contact with the end user for further up-sales, and maintenance services that could add up to significantly larger revenues over time than a simple product delivery and sale.

Additionally, we will see a shift in IoT from technology- and vendor-specific solutions to agnostic solutions. The primary driver for this change is to help reduce the complexity of incorporating and integrating point solutions that slow down the number of purchases but increases the value of the solutions services provided to the customer.

In terms of defining standards, the good news is many organizations already are taking leadership roles.

Here are just a few groups that are furthering education, standards, and best practices for creating, integrating, deploying, and maintaining an IoT program:

- Standards development efforts, like OneM2M,³ set the direction for the technical requirements for IoT interoperability and architecture. As other industries adopt IoT, they will add to their own unique standards the requirements for data and analytics that IoT benefits can provide.
- The IoT M2M Council (IMC)⁴ is focused on proving the business case of IoT technology to customers who adopt it. It aims to stand for IoT applications and connectivity as its own global industry and not viewed through the narrow lens of a technical standard or single vertical industry.
- The Healthcare Information and Management Systems Society (HIMSS)⁵ is a healthcare-focused organization that leads global endeavors to optimize health engagements and care outcomes through information technology. Its large membership is representative of the healthcare industry and the organization is helping shape interoperability standards that will be successfully adopted.



³ www.onem2m.org

⁴ www.iotm2mcouncil.org

⁵ www.himss.org

Beyond the consolidation in the industry that is typical of any fast-growing, dynamic sector, IoT companies will move to providing fully enabled solutions instead of just products and basic services.

SECURITY CONCERNS WILL CONTINUE

Security concerns could slow down IoT adoption if they aren't approached thoughtfully. Given the massive amount of data generated from information-intensive IoT applications, the ever-increasing number of connected devices, and the need to provide secure connectivity, enterprises must be careful not to risk the privacy and security of individuals.

Although defending against all sophisticated cyber-attacks is not entirely possible for end users of IoT technologies on their own, following industry-proven best practices and designing programs with security and privacy in mind is vital to the successful adoption of IoT. Data security will become a very significant part of the IoT budget for most businesses.

Additionally, if an enterprise keeps security in mind from the initial phase of any product development effort ("Security by Design") and incorporates it at a level that is appropriate to the use case, then most concerns will be mitigated. Security implementations must be relative to the IoT application requirements, and must be affordable, scalable, and user friendly. With better insight, businesses may find that not all IoT use cases require the most robust of security measures.

For example, fitness trackers typically track activity such as the number of steps, number of calories burned, etc., so the devices may not need to activate or to be physically chained down. However, the complex, connected machinery that builds the parts of an airplane warrants physical security in addition to a series of virtual locks to prevent unauthorized access, as well as vetting of individuals who are permitted to remotely operate that machinery.

And finally, the enterprise should look at security as an ongoing process. Security should be about incremental changes and not pose a fundamental challenge that is insurmountable for business. Don't try to build security that is bulletproof for the next 20 or 30 years. To aim for such a goal is impossible, as we cannot anticipate what innovations will occur or how the market will change in the future.

OVER-THE-AIR UPDATES WILL BECOME THE NORM

As products incorporate more and more embedded software, the need for remote updates without requiring the attention of skilled technicians will be the new norm, since consumers cannot be expected to reliably perform manual updates.

Using over-the-air (OTA) updates to refresh code inside the product will be increasingly critical—to add new features, fix product bugs, and repair firmware that could be sources of security breaches that impact the networks or application.

With the large number of IoT devices inherent in many applications, the cost of technician dispatches to fix bugs that impact other devices or data networks would be prohibitive. OTA update capability must become the future norm in all deployed IoT units and applications. The added up-front cost to support these OTA updates easily will be repaid many times over, at the first occurrence of a need to update—for example, to fix a security breach due to cyber-attacks.



PRIVACY CONCERNS AND GOVERNMENT REGULATIONS

As data and information from IoT devices increases, the need to protect individuals and companies from privacy breaches will become paramount. In our complex world, the ability for consumers to understand the implications of privacy breaches is decreasing rapidly.

Enterprises deploying IoT applications must accept an essential obligation to protect the personal information of consumers to avoid problems such as identity theft and financial loss due to security breaches.



Government regulations to manage privacy concerns and enforce consistent compliance already are being legislated. For example, the European Union (EU) enacted the General Data Protection Regulation (GDPR) on May 25, 2018 to unify data privacy requirements across the entire geo-political landscape. The use of information related to Euro-citizens, companies, and employees comes with very specific expectations of privacy, and harsh penalties for serious violations that, potentially, could cause companies to go out of business if they are not compliant to the requirements of GDPR.

Other countries are exploring similar and related regulations, and companies deploying IoT solutions must follow these national laws and regulations to do business in those countries. For example, in many regions, data—regardless of whether it is privacy related or not—must not cross “national boundaries”. Thus, all processing of that data, including the business outcomes of that IoT solution, must be processed by systems local to that country.

IoT VALUE REALIZED THROUGH DATA ANALYTICS

As billions of devices, things, and processes become interconnected, they will create a massive volume of data that drives the need for IoT analytics to automatically deal with the deluge. IDC predicts that 25 million applications will be created, and 50 trillion gigabytes of data will be generated, by 2020.

Multiple analysts predict that by 2019, IoT-created data will be stored, processed, analyzed, and acted upon close to, or at the edge of the network, thus relieving part of the data proliferation challenge.

The need for data analytics will be so great that it will drive demand for jobs for individuals that specialize in data science. While the issue of how to get data off devices and into back-end systems is increasingly being resolved through edge computing, the challenge of who gets to monetize the data is not yet resolved. Also, there isn't clarity in who plays the role of the broker of all that data and who manages the data repository. There might be multiple roles or one. As the increasing demand for data scientists to help make sense of the data is met, organizations will start to utilize the data for predictive purposes so it drives better, more efficient organizational decisions rather than as a passive activity of analyzing the data after the fact.

THE FUTURE IS NOW

2017 was the year where the number of connected devices surpassed the number of humans living on the planet. And, as with any new, fast-growing technology, the Internet of Things is not without challenges. However, the opportunities are vast, and those that navigate the waters thoughtfully can find great success.



DIRECTORY OF IoT TERMS

168

ACRONYMS

178

GLOSSARY

e

d

c



ACRONYMS

#s

3GPP	3rd Generation Partnership Project
------	------------------------------------

2G	Second generation (as in cellular technology)
----	---

3G	Third generation (as in cellular technology)
----	--

4G	Fourth generation (as in cellular technology)
----	---

5G	Fifth generation (as in cellular technology)
----	--

A

ADAS	Advanced Driver Assistance Systems
------	------------------------------------

AES	Advanced Encryption Standard
-----	------------------------------

AMPS	Advanced Mobile Phone System
------	------------------------------

AMPQ	Advance Message Queuing Protocol
------	----------------------------------

ANSI-41	American National Standards Institute Standard 41 (control signal messaging on SS7)
---------	---

ANSI-95	American National Standards Institute Standard 95 (for CDMA cellular)
---------	---

ANSI-136	American National Standards Institute Standard 136 (for TDMA cellular)
----------	--

ANSI-2000	American National Standards Institute Standard 2000 (for CDMA 200 cellular)
-----------	---

API	Application Programming Interface
-----	-----------------------------------

ARP	Address Resolution Protocol
-----	-----------------------------

AWS	Amazon Web Services
-----	---------------------

B

BAN	Body Area Network
-----	-------------------

BLE	Bluetooth Low Energy
-----	----------------------

BSC	Base Station Controller
-----	-------------------------

BTS	Base Transceiver Station
-----	--------------------------

BYOC	Bring Your Own Carrier
------	------------------------

C

CAN	Controller Area Network
-----	-------------------------

CDMA	Code Division Multiple Access
------	-------------------------------

CoAP	Constrained Application Protocol
------	----------------------------------

D

DG	Distributed Generation
----	------------------------

DL or D/L	Download data to end node from server / target address
-----------	--

DoF	Degrees of Freedom
-----	--------------------

DR	Demand Response
----	-----------------

E

EAN	European Article Number
------------	-------------------------

ECU	Electronic Control Unit
------------	-------------------------

EDGE	Enhanced Data rates for GSM Evolution
-------------	---------------------------------------

EPC	Electronic Product Code
------------	-------------------------

ESD	Electrostatic Discharge
------------	-------------------------

ESN	Electronic Serial Number
------------	--------------------------

EV-DO	Enhanced Voice Data Only (or Enhanced Voice Data Optimized)
--------------	---

F

FAKRA	Fachnormenausschuss Kraftfahrzeugindustrie
--------------	--

FDMA	Frequency Division Multiple Access
-------------	------------------------------------

FOTA	Firmware Over-the-Air
-------------	-----------------------

G

GGSN	Gateway GPRS Support Node
-------------	---------------------------

GIS	Geographic Information System
------------	-------------------------------

GLONASS	Russian global navigation system
----------------	----------------------------------

GNSS	Global Navigation Satellite System
-------------	------------------------------------

GPRS	General Packet Radio Service
-------------	------------------------------

GPS	U.S. Global Positioning System
------------	--------------------------------

GSM	Global System for Mobile communications
------------	---

H	
----------	--

HDFS	Hadoop Distributed File System
-------------	--------------------------------

HEM	Home Energy Management
------------	------------------------

HEMS	Home Energy Management System
-------------	-------------------------------

HetNet	Heterogeneous Network
---------------	-----------------------

HLR	Home Location Register
------------	------------------------

HSDPA	High-Speed Downlink Packet Access
--------------	-----------------------------------

HSPA	High-Speed Packet Access
-------------	--------------------------

HSPA+	Enhanced or Evolved High-Speed Packet Access
--------------	--

HSUPA	High-Speed Uplink Packet Access
--------------	---------------------------------

HVAC	Heating, Ventilation, and Air Conditioning
-------------	--

I	
----------	--

I2C	Inter-Integrated Circuit
------------	--------------------------

IaaS	Infrastructure as a Service
-------------	-----------------------------

ICCID	Integrated Circuit Chip Identifier
--------------	------------------------------------

ICS	Industrial Control System
------------	---------------------------

ICT	Information and Communications Technologies
IETF	Internet Engineering Task Force
IIoT	Industrial Internet of Things
IMEI	International Mobile Equipment Identifier
IMS	Intelligent Multi-Media System
IMSI	International Mobile Subscriber Identifier
IoE	Internet of Everything
IoT	Internet of Things
IPSEC	Internet Protocol Security
IPv6	Internet Protocol Version 6
IS-136	Interim Standard 136
IS-95	Interim Standard 95
ISDN	Integrated Services Digital Network
ISM Bands	Industrial, Scientific, and Medical Bands
ITS	Intelligent Transportation System
ITU	International Telecommunications Union
IVI	In-Vehicle Infotainment

J

JSON	JavaScript Object Notation
-------------	----------------------------

L

L2TP	Layer 2 Tunneling Protocol
-------------	----------------------------

LAN	Local Area Network
------------	--------------------

LED	Light Emitting Diode
------------	----------------------

LPWA	Low Power Wide Area
-------------	---------------------

LPWAN	Low Power Wide Area Network
--------------	-----------------------------

LTE	Long-Term Evolution
------------	---------------------

LTE-M	LTE category M1
--------------	-----------------

M

M2M	Machine-to-Machine
------------	--------------------

MAC	Media Access Control
------------	----------------------

MCU	Micro-Controller Unit
------------	-----------------------

MDN	Mobile Directory Number
------------	-------------------------

MEID	Mobile Equipment Identifier
-------------	-----------------------------

MEMS	Micro-Electro-Mechanical Systems
-------------	----------------------------------

MIMO	Multiple Input, Multiple Output (in regards to antennas)
-------------	--

MMS	Multimedia Messaging Service
MNO	Mobile Network Operator
MQTT	Message Queue Telemetry Transport
MS	Mobile Station
MSC	Mobile Switching Center
MSISDN	Mobile Station ISDN
MVNO	Mobile Virtual Network Operator
N	
NB-IoT	Narrowband IoT
NFC	Near Field Communication
P	
PaaS	Platform as a Service
PAN	Personal Area Network
PCB	Printed Circuit Board
PDU	Power Distribution Unit
PERS	Personal Emergency Response System
PoE	Power over Ethernet
PPTP	Point-to-Point Tunneling Protocol

PRL	Preferred Roaming List
-----	------------------------

PXE	Preboot Execution Environment
-----	-------------------------------

Q

QoS	Quality of Service
-----	--------------------

R

RADIUS	Remote Authentication Dial-In User Service
--------	--

REST	Representational State Transfer
------	---------------------------------

RF	Radio Frequency
----	-----------------

RFC	Request for Comment
-----	---------------------

RFID	Radio Frequency Identification
------	--------------------------------

RPMA	Random Phase Multiple Access
------	------------------------------

S

SaaS	Software as a Service
------	-----------------------

SBC	Single Board Computer
-----	-----------------------

SCADA	Supervisory Control and Data Acquisition
-------	--

SDN	Software-Defined Network
-----	--------------------------

SDO	Standards Development Organization
-----	------------------------------------

SGSN	Serving GPRS Support Node (also see GGSN)
------	---

SIM	Subscriber Identity Module
------------	----------------------------

SMA	Sub-Miniature version A
------------	-------------------------

SMS	Short Message Service
------------	-----------------------

SMSC	Short Message Service Center
-------------	------------------------------

SOAP	Simple Object Access Protocol
-------------	-------------------------------

SoC	System on a Chip
------------	------------------

SS7	Signaling System 7
------------	--------------------

STOMP	Simple (or Streaming) Text-Oriented Message Protocol
--------------	--

T	
----------	--

TCP / IP	Transmission Control Protocol / Internet Protocol
-----------------	---

TDMA	Time Division Multiple Access
-------------	-------------------------------

TETRA	Terrestrial Trunked Radio
--------------	---------------------------

U	
----------	--

UART	Universal Asynchronous Receiver / Transmitter
-------------	---

UBI	Usage-Based Insurance
------------	-----------------------

UL or U/L	Uplink or Upload
------------------	------------------

UMTS	Universal Mobile Telecommunications System
-------------	--

URI	Uniform Resource Identifier
-----	-----------------------------

URL	Uniform Resource Locator
-----	--------------------------

V

V2I	Vehicle-to-Infrastructure
-----	---------------------------

V2V	Vehicle-to-Vehicle
-----	--------------------

V2X	Shorthand for combining vehicle-to-vehicle (V2V), vehicle-to-infrastructure (V2I), and vehicle to anything
-----	--

VLR	Visitor Location Register
-----	---------------------------

VPN	Virtual Private Network
-----	-------------------------

W

WAN	Wide Area Network
-----	-------------------

WAP	Wireless Application Protocol
-----	-------------------------------

WAVE	Wireless Access in Vehicular Environments
------	---

Wi-Fi	Wireless Fidelity
-------	-------------------

GLOSSARY

#s

1xEV-Do	1 times Evolution Data Optimized (used in ANSI-2000 CDMA).
1xRTT	1 times Radio Transmission Technology (used in ANSI-2000 CDMA).
2.4 GHz	Wireless band commonly used in technologies such as Wi-Fi, Bluetooth, and ZigBee. This unlicensed band also is used by some LPWA technologies.
2G	Second-generation cellular technology. This technology is starting to be sunset in U.S. and elsewhere.
3G	Third-generation cellular technology offering improved data transfer rates over 2G. Increased capacity and data speeds with additional protocols.
3GPP	3rd Generation Partnership Project (3GPP) is a collaborative project aimed at developing globally acceptable specifications for third-generation (3G) mobile systems (GSM).
3GPP2	3rd Generation Partnership Project (3GPP2) is a collaborative project aimed at developing globally acceptable specifications for third-generation (3G) mobile systems (CDMA).
4G	Fourth-generation cellular technology and the latest upgrade to the GSM network, providing greater data transfer speeds. Also known as LTE.
5G	Fifth-generation cellular technology.
6LoWPAN	Communication protocol that compresses Ipv6 packages for small, low-power devices so as to let them communicate within the IoT.

802.11ah	Wi-Fi protocol that uses sub-1 GHz license-exempt bands (as opposed to conventional Wi-Fi that operates in the 2.4 GHz and 5 GHz bands).
868 MHz	License-free RF band mostly used for short-range applications, such as thermostats, security alarms, and industrial uses.
92 MHz	License-free RF band used for short-range applications. The low frequency allows for better penetration through walls and obstacles. Has a low data transfer rate.
A	
Acceleration Sensing	A MEMS concept referring to the increase in movement of an object from one point to another along a straight line or axis. Applications include remote control, pointing devices, gesture recognition, fitness monitoring equipment.
Accelerometer	Tool that measures changes in acceleration in the unit in which it is installed. Used to measure acceleration, tilt, vibrations.
Access Control	A system that determines who, when, and where people can enter or exit a facility or area. Used for electrical systems, wireless locks, cybersecurity, etc.
Access Point	A Wi-Fi node that allows users entry to a network, typically a LAN.
Active Sensor	A sensing device that requires an external source of power to operate.
Actuator	A device that introduces motion by converting electrical energy into mechanical energy in an electromechanical system. An actuator also may stop motion by clamping or locking.
Address Resolution Protocol (ARP)	Communication protocol used to convert an IP address into a physical address. This way, computers can communicate with each other, despite only knowing each other's IP addresses, by sending an ARP request that informs them about the other computer's MAC address.

Advanced Driver Assistance Systems (ADAS)	Digital features incorporated into vehicles to enhance driver safety. ADAS functionality includes digital vision for lane departure warnings, blind spot detection, radar for collision avoidance, and V2V communications for multiple vehicles operating near each other.
Advanced Message Queuing Protocol (AMQP)	An open-source standard for business messaging communications. Main features include message orientation, queuing, routing, reliability, security.
Advanced Metering Infrastructure	Architecture for automated, two-way communications between a smart utility meter with an IP address and a utility company.
Advanced Mobile Phone System (AMPS)	An analog cellular mobile system using FDMA. Analog AMPS has been supplanted by digital.
Amazon Web Services (AWS)	The name given to a collection of remote computing services, offered by Amazon.com, that combine to make a cloud computing platform.
Anomaly Detection	Statistical technique that determines the patterns that are normal and then identifies items that do not conform to those patterns. Unlike simple classification where classes are known in advance, in anomaly detection, the users don't know what they are looking for in the data.
Application Programming Interface (API)	A collection of commands and protocols used to interact with an operating system, device, or software component. In IoT, an API lets the developer access the functionality of a device or sensors.
Application Software	Programs that enable specific end-user actions. The software uses the given potential provided by computers to form an application.
Arduino	A single-board micro-controller used for prototyping without having to deal with breadboards or soldering. Software is free and open source.

AT Commands	Attention commands are used to set data connections. The set of short string commands allow developers to set up calls with a modem, as well as perform far more complex tasks.
Audio Profile	Hardware profile used with Bluetooth applications that include custom AT commands and functionality dedicated to wireless streaming of audio. Examples include A2DP, which allows for streaming of audio to devices such as speakers, where an audio gateway profile allows for two-way audio communications used in devices such as headsets.
Augmented Entity	A physical entity is represented by a virtual entity on the digital level. An augmented entity combines the two and stands for any combination of the two entities.
B	
Band	A range of frequencies used by a technology for communications purposes. For example, the 2.4 GHz band is used for Wi-Fi and Bluetooth communications.
Bandwidth	In signal processing, the measure of the width of a range of frequencies.
Base Station	The radios and other equipment at the cell sites that are used to communicate with cellular devices.
Beacons	Low-cost devices that communicate with smartphone apps indoors. Beacons use Bluetooth and are key enablers for the smart retail category, triggering messages as consumers pass through locations or near products.
Big Data	Data sets so large that they cannot be used with traditional database tools. Big Data often requires massively parallel computing resources to access, curate, and analyze.

Bluetooth	Short-range wireless technology standard that operates on the 2.4 GHz band. Bluetooth can be used for sending both data and audio. Bluetooth devices can be set up with different hardware profiles to help perform specific tasks, such as with an audio adapter, an audio headset, or keyboard profiles.
Bluetooth LE	Bluetooth Low Energy (also known as Bluetooth 4.0). Offers lower power usage for devices.
Body Area Network (BAN)	A wireless network of wearable computing devices and sensors, which may be embedded inside the body. A BAN also may be called a WBAN, as in wireless body area network. Key use case is for healthcare applications.
Bring Your Own Device (BYOD)	Enterprise term recognizing that people are bringing their own Wi-Fi-enabled devices into the corporate network.
Broadband	A high-speed, always-on data communications channel.
Brownfeild	Describes the problem and the process of having to consider already existing systems when implementing new software systems.
Business Logic	Describes the processes that are necessary to enable or execute communication between an end user and a database or server. These processes describe how data is transmitted, transformed, or calculated.
C	
CAN Bus	A message-based, multi-master serial protocol for transmitting and receiving vehicle data within a CAN. The CAN Bus connects multiple Electronic Control Units (ECUs), also known as nodes. Initially designed for automotive, the CAN Bus can be adapted to aerospace, commercial vehicles, industrial automation, and medical equipment.

Carrier	A company that provides telecommunications services.
Cellular Modem	Allows a device to access the Internet using cellular mobile networks. Devices can be configured to remotely connect to a server or device to enable off-site access and data collection.
Cellular Router	Allows connected devices to access servers and other devices by making an IP connection through the cellular mobile network. Routers allow for multiple devices to be connected and controlled, while offering extra device and data transfer security.
Chief IoT Officer (CIoT)	The CIoT coordinates the integration of IoT into the enterprise. Successful CIoTs will break down silos between disciplines, such as big data, data analytics, security, communications protocols, etc.
Class 1 Bluetooth	Offers a greater wireless data transfer distance (more than 100m, up to 1km) with greater power consumption (100mW).
Class 2 Bluetooth	Short-range wireless data transmission (10-20m), which has low power consumption of about 2.5mW.
Cloud	Cloud computing is an information technology paradigm that enables ubiquitous access to shared pools of configurable system resources and higher-level services that can be provisioned rapidly with minimal management effort, via the internet. Cloud computing relies on sharing of resources to achieve coherence and economy of scale.
Cloud Communications	Communication services provided by third parties that can be accessed and used through the internet.
Cloud Computing	An approach where information technology capacities (such as storage or applications) are separated from the individual computer and are supplied through the internet at the user's demand. The "as-a-Service" moniker is sometimes used for cloud computing services, such as Software-as-a-Service, Platform-as-a-Service, and Infrastructure-as-a-Service. The backend for many IoT devices may be delivered via the cloud.

Code Division Multiple Access (CDMA)	Communications method used by many cell phone companies. CDMA is an example of multiple access, where several transmitters can send information simultaneously over a single communication channel. This allows several users to share a band of frequencies (see bandwidth). To permit this without interference between users, CDMA employs spread spectrum technology and a special coding scheme (where each transmitter is assigned a code).
Communication Model	Communication models try to capture, explain, simplify, and then model communication. One of the oldest and most famous models, the Shannon and Weaver Model, was created in 1949.
Companion Device	In wearables, a companion device requires a parent device, such as a smartphone, to fully operate. The opposite would be a standalone device that can do everything on its own.
Connected Healthcare	Connected health encompasses all advancements in the medical industry that relate to IoT communication and remote sensing.
Connected Home	A connected home is networked to enable the interconnection and interoperability of multiple devices, services, and apps, ranging from communications and entertainment to healthcare, security, and home automation. These services and apps are delivered over multiple interlinked and integrated devices, sensors, tools, and platforms. Related to Smart Home.
Constrained Application Protocol (CoAP)	This software protocol is used in small electronics devices and serves as the data encoding protocol between those devices.
Controller Area Network (CAN)	A controller area network (CAN) is a serial bus network of microcontrollers that connects devices, sensors, and actuators in a system or sub-system for real-time control applications. In automobiles, a CAN connects Electronic Control Units (ECUs) using a multi-master serial bus (the CAN bus) to control actuators or receive feedback from sensors. ECUs can be sub-systems, such as airbags, transmission, antilock brakes, or engine control. The standard consists of ISO 11898-1 and ISO 11898-2.

Control Network	A network of nodes that collectively monitors, senses, and controls or enables control of an environment for a specific purpose. A home appliance network is a one example of a control network.
------------------------	--

D	
----------	--

Dashboard	A user interface that presents key data in a summarized form, often as graphs or other widgets. Derived from the classic automobile dashboard, the design of the interface depends on what data needs to be monitored or measured.
------------------	--

Data Center	A collective term for the physical site, network elements, systems, etc. that supports computing and network services.
--------------------	--

Data Lake	A data lake is a massive data repository, designed to hold raw data until it is needed and to retain data attributes so as not to preclude any future uses or analysis. The data lake is stored on relatively inexpensive hardware, and Hadoop can be used to manage the data, replacing OLAP as a means to answer specific questions. Sometimes referred to as an “enterprise data hub,” the data lake and its retention of native formats sits in contrast to the traditional data warehouse concept.
------------------	---

Degrees of Freedom (DoF)	A concept used in MEMS to describe the directions in which an object can move and the number of independent variables in a dynamic system.
---------------------------------	--

De-identification	The stripping away of personally identifiable information from data prior to its use. The process must include the removal of both direct identifiers (name, email address, etc.) and the proper handling of quasi-identifiers (sex, marital status, profession, postal code, etc.).
--------------------------	--

Demand Response	Demand response can reduce electrical price volatility during peak demand periods and help avoid system emergencies.
------------------------	--

Device Attack	An exploit that takes advantage of a vulnerable device to gain network access.
----------------------	--

DIN Rail	Metal rail used for mounting electrical equipment and racks
Distributed Generation (DG)	Decentralized, modular, and flexible power generation located close to the serviced loads. Distributed micro-grids can control smaller areas of demand with distributed generation and storage capacity.
DNP3 Protocol	An open, standards-based protocol for the electric utility industry with interoperability between substation computers, remote terminal units, intelligent electronic devices, and master stations. Groups of enabled things are organized into namespaces.
Domain Model	A model that contains all areas and terms related to a certain field of interest. It includes attributes, relations, constrains, acts, etc., that are relevant for a certain task.
Downlink (DL or D/L)	The process of downloading data onto an end node from a server / target address. In a cellular network, this would be seen as data being sent from a cellular base station to a mobile handset.
E	
eHealth	Telemedicine, telehealth. Related to mHealth. Medical processes and applications through information and computer technologies.
Electronic Control Unit (ECU)	Also known as a node, an Electronic Control Unit is a device, such as a sensor or actuator, that is connected to other devices via a CAN Bus. A vehicle can contain dozens of ECUs for functions such as mirror adjustment, window power, airbags, cruise control, entertainment, and, most significantly, engine control. To form a CAN, two or more ECUs are needed.
Electronic Serial Number (ESN)	Unique identification numeral for mobile devices in CDMA. Replaced by the MEID.
Electrostatic Discharge (ESD)	This sudden flow of electricity can occur if two electrical objects, with different electrical charges, come in contact with each other. The difference in charge often is due to friction. Sometimes, the short process is accompanied by sparks, as can be seen with lightning. ESD can lead to severe damage to electrical devices.

Embedded Firmware	The flash memory chip that stores specialized software running in a chip in an embedded device to control its functions. (Firmware = software for hardware.)
Embedded System Security	The reduction of vulnerabilities and protection against threats in software running on embedded devices.
Energy-Harvesting Technologies	Technologies that use small amounts of energy from close proximity to power small wireless devices. Applications can be found in wireless sensor networks or wearable tech. Energy sources include sun, wind, or kinetic energy.
Enhanced Data Rates for GSM Evolution (EDGE)	An enhancement made to 2G GSM networks to improve data transfer speeds and provide downlink speeds of up to 1 Mbit/s and uplink speeds of up to 400 Kbit/s. It builds on available GSM or GPRS standards and is integrated easily into existing networks.
Enhanced Voice Data Only (EV-DO)	Enhanced Voice Data Only (or Enhanced Voice Data Optimized).
EPCglobal	Joint venture set up to achieve worldwide adoption and standardization of Electronic Product Code (EPC) technology.
EtherCAT	A fieldbus system that allows for real-time Ethernet. It helps to achieve short data update times, accurate synchronization, and low hardware costs, so it can be used specifically for automated or control systems. (CAT stands for Controller and Automation Technology.)
European Article Number (EAN)	Used to mark and identify products. Since 2009, it also is called Global Trade Item Number (GTIN). The number usually is found beneath barcodes and consists of up to 13 digits (EAN 13 barcode).

F

Fachnormenausschuss Kraftfahrzeugindustrie (FAKRA)	This is a type of SMB connector used in the automotive industry for connecting coaxial RF connectors.
Fast Data	This is the application of Big Data analytics to smaller data sets in near real time or in real time to solve a problem or create business value.
Firmware	Programming that is written to the read-only memory (ROM) of a computing device. Firmware, which is added at the time of manufacturing, is used to run user programs on the device.
Firmware Over-the-Air (FOTA)	The process of updating an operating system and software over the network, rather than having the consumer come into a service center for updates.
Fitness Band	Activity tracker worn on the wrist, with sensors specifically related to exercise and activity measuring. In contrast to a smartwatch that may include fitness / activity tracking features, a fitness band is dedicated to fitness.
Fleet Management	A broad term referencing a range of solutions for vehicle-related applications. A fleet management solution typically is a vehicle-based system that incorporates data logging, satellite positioning, and data communication to a back-office application.
Fog Computing	A distributed computing infrastructure in which some application services are handled at the network edge in a smart device and some application services are handled in a remote data center—in the cloud.
Form Factor	The physical size, pin-out, and configuration of a component. A family range of modules, for example, may include 2G, 3G, and 4G variants to allow PCB designers to design in one module but allow for future upgrades through the product family's road map.
Frequency Division Multiple Access (FDMA)	The division of the frequency band allocated for wireless cellular telephone communication into channels, each of which can carry a voice conversation or digital data.

G

Gateway	A link between two computer systems or programs, which allows them to share information with each other. The router for your home internet is one type of gateway.
Gateway GPRS Support Node (GGSN)	A main component of a GPRS network that supports the networking between the GPRS network and external packet-switched networks. See also SGSN.
General Packet Radio Service (GPRS)	A wireless communications standard on 2G and 3G cellular networks, which supports a number of bandwidths and provides theoretical data rates of 56-114 kbps. As cellular companies move to more advanced networks, GPRS networks may be more cost-effective for IoT networks.
Geofence	A virtual border applied to a physical space. For example, geofencing might be defined around a nursery, and when a mobile device crosses the nursery boundary, an alert is generated. Geofences may be dynamically created and, in a telematics application, can encompass entire neighborhoods or cities.
Geographic Information System (GIS)	The combination of hardware, software, and data that captures, manages, analyzes, and presents many kinds of geographic data. GIS and location intelligence applications can be the foundation for location-enabled services.
Global Navigation Satellite System (GNSS)	General term for the multiple constellations of satellite navigation systems.
Global Positioning System (GPS)	A U.S. system of satellites and radio transmissions that is used to locate GPS-enabled hardware anywhere on the planet, with a very high degree of accuracy.
Global System for Mobile Communication (GSM)	The most widely used digital cellular network and the basis for mobile communications, such as phone calls and short message services (SMS).
Greenfield	In contradiction to brownfield, a greenfield project is one where no consideration of previous systems is needed.

H

Hadoop	A Java-based distributed programming framework for processing large data sets. An application can be broken down into numerous small parts, called fragments or blocks, that can be run on any node in the cluster. Hadoop is part of the Apache Project, sponsored by the Apache Software Foundation.
Hadoop Distributed File System (HDFS)	The primary distributed storage used by Hadoop applications. A HDFS cluster has a NameNode that manages the file system metadata and DataNodes to store the actual data.
Handoff	The transfer of a wireless call in progress from one transmission site to another site without disconnection.
Haptic Technology or Haptics	Haptic technology (Haptics or touch feedback) applies tactile sensations to human interactions with machines. The simplest example is the actuator that vibrates a cell phone, but more advanced haptics can detect the pressure applied to a sensor, affecting the response.
Heating, Ventilation, and Air Conditioning (HVAC)	These systems cover both vehicular and indoor building comfort control.
Heterogeneous Network (HetNet)	Small cell networks using both macro and small cells. HetNets allow mobile operators to better utilize their data networks' capacity.
High-Speed Downlink Packet Access (HSDPA)	Increases the capacity of UMTS / 3G bandwidth to allow for faster download speeds for connected devices.
High-Speed Packet Access (HSPA)	An improvement made to data speeds over 3G technology through the addition of two new protocols; HSDPA and HSUPA. It offers potential downlink speeds of 14 Mbit/s and downlink of 5.76 Mbit/s.
High-Speed Uplink Packet Access (HSUPA)	An improvement made to UMTS to enable faster uploading of data from devices.

Home Automation	The automation of certain activities within a household. This can include automated control of lights, doors, and air conditioning, for example.
Home Energy Management (HEM)	Home Energy Management refers to technology that helps homeowners improve home energy efficiency while also giving them access to household products, such as tablets, smartphones, and computers.
Home Energy Management Systems (HEMS)	Any hardware and / or software system that can monitor and provide feedback about a home's energy usage and / or enable advanced control of energy-using systems and devices in the home.
Home Location Register (HLR)	Main database of permanent subscriber information for a mobile network.
Host	Computers that provide (or host) certain services or resources within a network that other participants within the network can then access and use. Hosts are the hardware basis for servers, as servers are run on hosts. Often, hosts are the central point in a company's data processing process.
Hybrid Cloud	A mix of public and private cloud. The distribution of services through private or public channels is decided upon by the users.
I	
iBeacon	A technology introduced by Apple that uses sensors to locate iOS or Android devices indoors and can send them notifications via Bluetooth Low Energy (BLE). This also can be used in stores or museums to give further information about nearby items.
Identifier	Also just ID, this marks objects for clear identification. Identifiers usually are letters, words, symbols, or numbers that can be used to create a code that reveals a definite identity after it is decoded.

Identity	Recognizable attributes that are linked to an object, a person, etc. Those attributes expose the entity and allow for clear identification. If two things have the exact same attributes, they usually have the same identity, and they can't be distinguished from each other.
Industrial Control System (ICS)	Computer hardware and software that monitor and control industrial processes that exist in the physical world, where operator-driven supervisory commands can be pushed to remote station devices. Industries such as electrical, water, oil, and gas are typical ICS users.
Industrial Internet of Things (IIoT)	A sub-discipline of IoT, encompassing connected large-scale machinery and industrial systems, such as factory-floor monitoring, HVAC, smart lighting, and security. Equipment can send real-time information to an application so operators can better understand how efficiently that equipment is running. Also referred to as Industrial IoT.
Industrial, Scientific, and Medical (ISM) Bands	An unlicensed part of the RF spectrum used for general purpose data communications. In the U.S., the ISM bands are 915 MHz, 2.4 GHz, and 5.5 GHz, whereas 2.4 GHz is the global unlicensed frequency, which has increasing amounts of interference.
Industrie 4.0	Invoking a fourth Industrial Revolution, Industrie 4.0 creates intelligent manufacturing networks where decentralized smart factories can communicate and react to each other autonomously. The term, also known as Industry 4.0, was first used at the Hannover Messe in 2011.
Industry 4.0	Industry 4.0 is a project introduced by the federal government of Germany and refers to the fourth Industrial Revolution. It is a strategy which aims to make better use of current and future IT-capacities in traditional industries.
Inertial Measurement Unit (IMU)	A MEMS module that measures angular velocity and linear acceleration using an accelerometer triad and an angular rate sensor triad. Other IMU sensors may include magnetometers and pressure sensors.

<p>Information and Communication Technologies (ICT)</p>	<p>The ICT industry provides access to information through telecommunications. The communications technologies can be items, such as the internet, VOIP, wireless networks, or mobile phones.</p>
<p>Infrastructure as a Service (IaaS)</p>	<p>An on-demand business model for IT capacities. Instead of owning IT-infrastructure or server space, you rent and pay for it on a per-use basis. Those capacities usually are owned, maintained, and provided by a cloud service.</p>
<p>Insurance Telematics</p>	<p>Vehicular tracking devices used by automobile insurance companies to alter rates based on driver behavior. Tracks a multitude of driver / driving related items, including hard braking, mileage, speed, hard turns, and much more.</p>
<p>Intelligent Device</p>	<p>Any type of equipment, instrument, or machine that has its own computing capability. As computing technology becomes more advanced and less expensive, it can be built into an increasing diversity of devices. The list of uses includes personal and handheld computers, cars, medical instruments, geological equipment, home appliances, and more.</p>
<p>Intelligent Multimedia Systems (IMS)</p>	<p>Fundamental tools in the retrieval and dissemination of data, as well as to enable face-to-face interaction across different geographies.</p>
<p>Intelligent Transportation System (ITS)</p>	<p>An application of advanced information and communications technology for surface transportation enabling enhanced safety and mobility while reducing environmental impacts.</p>
<p>Inter-Integrated Circuit (I2C)</p>	<p>I2C, pronounced I-squared-C, is a serial bus that provides communication between sensors and microcontrollers, such as the Arduino.</p>
<p>International Mobile Equipment Identifier (IMEI)</p>	<p>The unique number used in GSM to identify mobile devices on individual operator networks—conceptually similar to the MEID in CDMA.</p>

International Mobile Subscriber Identifier (IMSI)	The unique number used in GSM, CDMA, and LTE to identify SIM cards on their individual operator networks.
International Telecommunications Union (ITU)	A specialized agency of the United Nations responsible for issues concerning information and communication technologies.
Internet of Everything (IoE)	The IoE amasses the technologies found in M2M and the IoT and expands them with an even greater accumulation of data and inferences. IoE is defined as a networked connection of four key elements, including: People (social networks, health and fitness sensors, and more); Things (physical sensors, measuring devices, actuators, and other items generating or receiving data); Processes (leveraging connectivity among data, things, and people to add value, as in the use of smart fitness devices and social networks to advertise healthcare offerings to prospective customers); and Data (raw data analyzed and processed into useful information to enable intelligent decisions and control mechanisms).
Internet of Things (IoT)	The Internet of Things moves beyond the scope of M2M, encompassing and surpassing it in functionality by adding devices and electronic equipment with embedding sensors, control systems, and processors that enable communication across a multi-node, open network of objects. The IoT includes any object, outfitted with sensors, that has the ability to gather and transfer data over a network.
Internet Protocol Security (IPSEC)	A set of protocols that provide authentication and encryption to Internet Protocol (IP) packets, adding an extra layer of security on IP communications.
Interoperability	The ability of two or more systems or components to work together and exchange and use information effectively.

In-Vehicle Infotainment (IVI)	Systems integrated into automobiles that deliver both entertainment and information content. Typical IVI applications include managing audio, listening to or sending SMS, making voice calls, navigating, and using rear-seat entertainment, as well as interfacing with smartphone-enabled content, such as traffic conditions, sports scores, and weather forecasts.
IoT Privacy	Internet of Things privacy. The special considerations required to protect the data of individuals from exposure in an IoT environment, where almost any physical or logical entity or object can be given a unique identifier and the ability to communicate autonomously over the internet or similar IP network.
IoT Security	Internet of Things security. The area concerned with safeguarding connected devices and networks in the Internet of Things.
IP Devices	All devices within a network that are labeled with an IP address.
IPv6	IP addresses serve to identify devices on the internet. IPv6 is the newest internet address format, which provides more addresses than the IPv4 address format.
IPv6 Address	A 128-bit alphanumeric string that identifies an endpoint device in the Internet Protocol Version 6 (IPv6) addressing scheme.
IRIDIUM	A satellite communication constellation that provides global voice and data coverage through its satellite network, operating on the 1618.85 to 1626.5 MHz frequencies.
J	
JavaScript Object Notation (JSON)	Used as a lightweight alternative to XML for organizing data, JSON is text-based and human-readable. The format uses “name : object” pairs to organize the data.

L

Layer 2 Tunneling Protocol (L2TP)	A tunneling protocol used to support virtual private networks (VPNs) or as part of the delivery of services by ISPs. It does not provide any encryption or confidentiality by itself, relying on an encryption protocol that it passes within the tunnel to provide privacy.
Light-Emitting Diode (LED)	A semiconductor that generates light via electroluminescence. Infrared LEDs can be used for the remote control units for many consumer electronics.
Link Budget	An accounting of all of the losses in a wireless communication system. In order to “close the link,” enough RF energy has to make it from the transmitter to the receiver.
Local Area Network (LAN)	A network of devices in relatively close proximity. The two most common communications technologies used in LANs are Ethernet and Wi-Fi.
Long-Term Evolution (LTE) / 4G	LTE, often referred to as 4G, is a cellular network type offering superior data transfer speeds than its predecessor, 3G. It is part of the GSM upgrade path. Portable devices now can access data at high broadband speeds through LTE.
LoRa	LoRa is a proprietary, chirp spread spectrum radio modulation technology for LPWAN. LoRa uses license-free sub-gigahertz radio frequency bands (such as 169 MHz, 433 MHz, 868 MHz (Europe) and 915 MHz (North America)).
LoRa Alliance	The LoRa Alliance is an open, non-profit organization dedicated to promoting the interoperability and standardization of low-power wide area network (LPWAN) technologies to drive implementation of the IoT. LPWANs enable IIoT, civic, and commercial applications. LPWANs have lower costs and power requirements and longer range than mobile networks. Those advantages mean they can enable a much wider range of IoT applications, which have been constrained by budgets and power issues.

LoRaWAN	LoRaWAN is a media access control layer protocol for managing communication between LPWAN gateways and end-node devices, maintained by the LoRa Alliance. LoRaWAN defines the communication protocol and system architecture for the network while the LoRa physical layer enables the long-range communication link.
Low-Power Wide Area (LPWA)	LPWA networks are designed for IoT applications that have low data transmission rates, need long battery lives, can provide low-cost services, sometimes operate in remote or hard to reach locations (underground or geographically dispersed), and be easy to deploy across basically every business sector, including manufacturing, automotive, energy, utilities, agriculture, healthcare, wearables (for humans or animals), or transport.
Low-Power Wireless Sensor Network	A group of spatially distributed, independent devices that collect data by measuring physical or environmental conditions with minimal power usage.
M	
Machine Authentication	The authorization of an automated human-to-machine or machine-to-machine communication through verification of a digital certificate or digital credentials. Unlike user authentication, the process does not involve any action on the part of a human.
Machine Data	Also known as machine-generated data, this is digital information created by the activity of computers, mobile phones, embedded systems, and other networked devices.
Machine-to-Machine (M2M)	A term describing technology that allows for one connected device to communicate and exchange information with another connected device or sensor—without the assistance of a human.
MapReduce	A parallel processing model for handling extremely large data sets. First, a map process runs to reduce a data set to key value pairs (in sequence), and then a second reduce process combines those pairs into a smaller set of tuples (ordered lists).

Media Access Control (MAC)	The “layer 2” in a network that allows the physical medium (radio waves or wire signals) to be organized to pass data back and forth. For low-rate data wireless applications, the MAC has many implications on performance.
Mesh Networking or Mesh Network Topology	An ad hoc, local area network infrastructure where nodes communicate directly with each other without the need to pass through a central structure, such as an ISP. The only way to shut down a mesh network is to eliminate every node. Their adaptability makes them ideal for IoT applications.
Message Broker	A middleware program that translates a message from the messaging protocol of the sender into the messaging protocol of the receiver. This makes it much easier for two applications to communicate.
Message-Oriented Middleware (MOM)	Middleware that allows for synchronous, as well as asynchronous (queue), messaging between distributed systems.
Message Queuing Telemetry Transport (MQTT)	An open, lightweight IoT communications protocol for the transfer of telemetry messages.
mHealth or Mobile Health	This is the practice of medicine using mobile devices, particularly physiological sensors. Sensors may be enabled to communicate with a user’s mobile phone in a Body Area Network (BAN) configuration.
MicroController Unit (MCU)	A full computer on a single chip. The chip contains a CPU, a clock, non-volatile memory for the program (ROM or flash), volatile memory for input and output (RAM), and an I/O control unit.
Micro-Electro-Mechanical Systems (MEMS)	Miniaturized mechanical and electro-mechanical elements, typically used for measurements, such as accelerometers and gyroscopes. Systems-on-a-chip technology is used to embed mechanical devices, such as fluid sensors, mirrors, actuators, pressure and temperature sensors, and vibration sensors, on to semiconductor chips.

Mobile Directory Number (MDN)	The number a user would dial to reach a specific mobile phone. Used in CDMA—conceptually similar to the MSISDN in GSM.
Mobile Equipment Identifier (MEID)	Unique identification numeral for mobile devices used in CDMA—conceptually similar to the IMEI in GSM and LTE..
Mobile Network Operator (MNO)	Companies that operate traditional mobile communications networks.
Mobile Station (MS)	A cellular radio handset or cellular IoT device.
Mobile Station International Subscriber Directory Number (MSISDN)	The telephone number to the SIM card in a mobile phone. Used in GSM—conceptually similar to the MDN in CDMA.
Mobile Switching Center (MSC)	The center of a network switching subsystem, associated with communications switching functions, routing SMS messages, and interfacing with other networks.
Mobile Virtual Network Operator (MVNO)	A wireless communications provider that leases the infrastructure over which it provides services.
Modbus	A communication protocol mainly used to connect electronic devices. The Modbus Master (for example, a computer) requests information from the Modbus Slaves (for example, electronic thermometers). Up to 247 Slaves can transmit data to one Master.
Multimedia Messaging Service (MMS)	A feature of mobile devices that allows transmission of images, video, or audio in addition to short text messages via standardized communications protocols. See Short Messaging Service (SMS).
Multiple-Input and Multiple-Output (MIMO)	A radio technology using multiple antennas at both the transmitter and receiver to improve communication performance.

N

**Narrowband-IoT
(NB-IoT)**

Narrowband-IoT is a LPWAN radio technology standard developed to enable a wide range of devices and services to be connected using cellular telecommunications bands. NB-IoT focuses specifically on indoor coverage, low cost, long battery life, and enabling a large number of connected devices.

**Near Field Communications
(NFC)**

Short-range wireless communication between devices, used in applications such as contactless mobile payments, transport ticketing, and phone-as-key.

O

On-Board Equipment (OBE)

Components of a Vehicle-to-Infrastructure (V2I) implementation located in a moving vehicle, communicating wirelessly with roadside equipment. OBE applications may interface with other vehicle systems via a CAN Bus.

Open Source

Software where the source code is available and can be modified and freely redistributed. Open source is the opposite of closed, proprietary systems.

Open VPN

Open-source software application that implements virtual private network (VPN) techniques for creating secure point-to-point or site-to-site connections in routed or bridged configurations and remote access facilities.

Operational Technology

As opposed to Information Technology (IT), this refers to technologies associated with control and automation.

P

Part 90 Licenses	Small parts of the RF spectrum that are made available in small areas to businesses for data or voice communications. Many smart grid providers use Part 90 licenses for wireless data.
Passive Sensor	A device that detects and responds to input from physical environments.
Penetration Testing, Pen Testing, or Pentest	A method of evaluating the security of a network or system from internal or external threats. This is part of a full security audit and typically exploits a combination of weaknesses to gain access and then evaluates the capability of the network's defenders to detect and respond to the penetration.
Personal Area Network (PAN)	Interconnected devices operating in the range of a single person, typically 10 meters. PANS are (mostly or exclusively) wireless, making the term basically indistinguishable from Wireless PANs (WPAN).
Personal Emergency Response System (PERS)	A mobile duress panic alarm component of a monitoring system, typically for the residential market. Modern PERS devices go beyond their origins as a mere push button to include MEMS and various other sensors.
Pervasive Computing	Another term for ubiquitous computing.
Physical Web	Google's open standard to allow IoT devices to communicate via web addresses. By using HTTP, users can walk up and access any smart device (such as parking meters and vending machines) without the overhead of dedicated mobile apps.
Platform as a Service (PaaS)	Platform as a Service is a category of cloud computing services that provides a platform allowing customers to develop, run, and manage applications without the complexity of building and maintaining the infrastructure typically associated with developing and launching an app.
Point-to-Point Tunneling Protocol (PPTP)	A method for implementing virtual private networks (VPNs).

Power Distribution Unit (PDU)	A physical device with multiple outlets that connects electrical power to recipient devices. PDUs can be simple, such as a mounted power strip, or more complex by having power filtering, UPS, load balancing, or intelligent monitoring incorporated in the device.
Power over Ethernet (PoE)	The capability to deliver enough power to operate a device over an Ethernet connection. PoE is useful in certain low-voltage applications, such as passive IP cameras.
Preboot Execution Environment (PXE)	The ability to manage power over a network connection. A PXE-enabled device can be shut down or restarted via a network connection.
Preferred Roaming List (PRL)	A database (especially in a CDMA-based wireless device) that tells how to find and connect to locally available wireless network(s). The function of the PRL is most important when a device is outside its home network and must seek out an alternate network.
Printed Circuit Board (PCB)	Mechanically supports and electrically connects electronic and electrical components using conductive tracks, pads, and other features etched from one or more sheet layers of copper laminated onto or between sheet layers of a non-conductive substrate.
Private Cloud	Information technology services supplied via the cloud but only within a single organization, for example, one company.
Public Cloud	Information technology services supplied via a cloud that are public and made available for everyone.
Python	A widely used open-source programming language that can be implemented in variety of ways, including in embedded applications. There is a large library base that can be used by Python applications, helping to minimize coding while speeding up development time.
Python Script Interpreter	A tool that lets you run Python code, something that now is being embedded directly into devices, such as cellular modules.

Q

Quality of Service (QoS)	Different services that regulate data transfer priorities to identify and control the quality with which a service can be accessed by users. This is especially important if a certain quality (for example, bandwidth) has to be guaranteed to ensure the functionality of a service.
---------------------------------	--

R

Radio Fingerprinting	An electronic process that identifies each individual wireless handset by examining its unique radio transmission characteristics.
Radio Frequency (RF)	Radio waves. This term generally means “wireless communication” when referred to in IoT discussions.
Radio Frequency Identification (RFID)	Generally speaking, this is the use of strong radio waves to “excite” enough current in a small tag to send a radio transmission back.
Random Phase Multiple Access (RPMA)	Random Phase Multiple Access LPWAN technology was designed from the bottom up to optimize wide area connectivity for IoT devices.
Raspberry Pi	Raspberry Pi is a series of small, single-board computers developed in the UK by the Raspberry Pi Foundation to promote the teaching of basic computer science in schools and in developing countries. The original model became far more popular than anticipated, selling outside its target market for uses such as robotics.
Remote Authentication Dial-In User Service (RADIUS)	A type of server responsible for receiving user connection requests, authenticating the user, and returning all configuration information necessary for the client to deliver service.
Remote Monitoring and Control	The increasingly automated monitoring and control of devices, technologies, or processes. Wireless devices that send information gathered directly to control centers often are used to achieve this.

Remote Sensing	The use of various technologies to make observations and measurements at a target that usually is at a distance or on a scale beyond those observable to the naked eye.
Representational State Transfer (REST)	An architecture for web standards, especially for the HTTP protocol. It simplifies design of network applications compared to, for example, SOAP.
RESTful Web Services	Web services that are realized within the REST architecture are called RESTful Web Services.
RF Sensitivity	The minimum magnitude of input signal you need based on a specified signal-to-noise ratio to achieve a minimum error rate.
RFID Tagging	A system using small radio frequency identification devices for tracking purposes. An RFID tagging system includes the tag itself, a read / write device, and a host application for data collection, processing, and transmission.
Roaming	Using a wireless device in an area outside its home coverage area.
S	
Sensor	A device used to measure a specific characteristic of the surrounding environment, such as temperature. The use of sensors and actuators to connect things to the physical world is a key component of IoT.
Sensor Analytics	Statistical analysis of data that is created by wired or wireless sensors.
Sensor Hub	A technology that connects sensor data and then processes the data. This way, the hub does part of a processor's data-processing job.
Serial Port Profile (SPP)	A hardware profile used with Bluetooth applications that includes custom AT commands and functionality dedicated to wireless data connections and serial cable replacement.

Shock Sensing	A MEMS concept referring to the detection of sudden impacts at a predetermined level. Typical applications include shut-off sensing, condition monitoring, and tap detection for data entry.
Short Message Service Center (SMSC)	The network element in a mobile telephone network that stores, forwards, converts, and delivers SMS messages.
Short Message Service (SMS)	A feature of mobile devices that allows transmission of short text messages via standardized communications protocols.
Sigfox	A low-bandwidth, proprietary, wireless protocol that offers excellent range and obstacle penetration for short messages, giving a new low-powered and cost-effective wireless transmission transport for IoT technologies.
Simple (or Streaming) Text Oriented Message Protocol (STOMP)	A protocol designed for working with message-oriented middleware, similar to HTTP. It allows clients to communicate with most of the message brokers, making it language agnostic.
Simple Object Access Protocol (SOAP)	A protocol specification for exchanging structured information in the implementation of web services in computer networks.
Single Board Computer (SBC)	A complete, functioning computer with all functions (I/O, processor, memory) located on one board. Popularized by the Raspberry Pi system, SBCs are constructed in direct contrast to traditional motherboards with plug-in cards for functions such as graphics and Ethernet.
Smart Buildings	Buildings that try to minimize costs and environmental impact. This is achieved by connected systems and efficient use of energy through new, automated technology that intelligently responds to certain aspects (available solar energy, temperature inside the building, etc.).

Smart Car An automobile that uses technology to support the driver and create a safer traffic environment. Different systems (inside and outside of car) are connected and communicate with each other to allow intelligent intervention in dangerous situations. Additional functionality includes (but not limited to) full stack telematics, as well as comprehensive global device, connectivity, and services delivery.

Smart Cities Smart cities are defined by more intelligent city infrastructure using modern information and communication technologies. Smart cities propose a more flexible adaptation to certain circumstances, more efficient use of resources, higher quality of life, more fluid transportation, and more. This is achieved via networking and integrated data exchange between humans and things.

Smart Grid A term referring to the application of networking capabilities and computer systems to the electric grid. For example, a smart grid would include smart meters at the point of delivery, allowing for real-time monitoring of usage and the adjustment of power settings on some appliances.

Smart Home The networking of household devices and systems through information and communication technology. This way, processes within a home can be monitored and controlled automatically to optimize quality of life, costs, security, and environmental impact. Related to Connected Home.

Smart Label A type of identification tag that contains more advanced technologies than conventional barcode data. Some common types of smart labels are QR codes, Electronic Article Surveillance (EAS) tags, and RFID tags.

Smart Meter An electronic device that measures and displays resource consumption (of water, gas, electricity, etc.) and communicates this information to the resource distributors and managers (such as utilities and municipalities) and even to consumers. This allows for a more efficient distribution, usage, pricing, and control of resources.

Software as a Service (SaaS)	A subscription-based model where a monthly fee is charged for using software, rather than an upfront purchase. SaaS and cloud computing can give cash-strapped enterprises and startups access to applications that might otherwise be too expensive to purchase outright.
Software-Defined Network (SDN)	An approach to networking that decouples control of information flow from the hardware and gives it to a software controller.
Standards Development Organization (SDO)	An organization whose primary activities are developing, coordinating, revising, amending, interpreting, or otherwise producing technical standards.
Structure Attenuation	The loss in intensity of radio waves through a medium (like radio waves through a brick wall).
Subscriber Identity Module (SIM)	A piece of hardware (the “smart card”) containing account information for a user on a cellular network. The SIM is inserted into a SIM holder in cellular devices.
Supervisory Control and Data Acquisition (SCADA)	An industrial control system typically used for geographically dispersed assets, often scattered over large distances. SCADA is applied to electrical utilities to monitor substations, transformers, and other electrical assets.
System on a Chip (SoC)	A single integrated-circuit technology that contains all the necessary circuits and parts for a complete system. A single microchip in a wearable device, for example, could contain an analog-to-digital converter, memory, logic control, I/O, etc.

T

Telematics	An IT concept regarding the long-distance transmission of data. In vehicles on the move, telematics refers to the integrated use of telecommunications and informatics, such as dashboard screens that show the vehicle's current position on a map or in centralized tracking applications. Telematics is an interdisciplinary field that can encompass telecommunications, vehicular technologies, road transportation, road safety, electrical engineering (sensors, instrumentation, wireless communications, etc.), and computer science (multimedia, internet, etc.).
Terrestrial Trunked Radio (TETRA)	This operates as a two-way transceiver and is used by emergency services, as well as on transport, such as rail and marine vessels. It operates on low frequencies split over four channels (ranging between 380 and 400 MHz for emergency services and higher for civilian use). The use of low frequencies allows for far greater transmission distances but lower data transfer rates.
Thread	A simplified IPv6-based mesh networking protocol geared to the smart home sector. Developed on low-cost 802.15.4 chipsets, Thread is designed for extremely low-power consumption.
Time Division Multiple Access (TDMA)	A channel access method for shared medium networks.
Transceiver	Short for transmitter-receiver. A transceiver both transmits and receives analog or digital signals. It normally is built into a network interface card.
Transmission Control Protocol / Internet Protocol (TCP / IP)	Core standard protocol for IP network communications.
Transponder	A wireless communications device that picks up and automatically responds to an incoming signal. The term is a combination of the words transmitter and responder. Transponders can be either passive or active.
TV Whitespace	A new FCC program that makes unused TV station bands available for temporary and controlled use in a small geographic area.

U

Ubiquitous Computing	The concept of embedding microprocessors in everyday things so they can communicate information continuously. Ubiquitous devices are expected to be connected constantly. Utility smart meters are an example of ubiquitous computing, replacing manual meter readers with devices that can report usage and modify power settings on machines, panels, etc.
Uniform Resource Identifier (URI)	The unique identifier makes content addressable on the internet by uniquely targeting items, such as text, video, images, and applications.
Uniform Resource Locator (URL)	A particular type of URI that targets web pages so that when a browser requests them, they can be found and served to users.
Universal Asynchronous Receiver / Transmitter (UART)	A microchip controlling a computer's interface to serial devices. It converts the bytes it receives from the computer along parallel circuits into a serial bit stream.
Universal Authentication	A network identity verification method that allows users to move from site to site securely without having to enter identifying information multiple times.
Universal Mobile Telecommunications System (UMTS)	Also referred to as 3G cellular technology, this is the third iteration of the GSM. It achieves improved data transfer speeds over 2G.
Uplink (UL or U/L)	This is the process of sending data from your device / computer to a server or target address. In a cellular network, this would be seen as data sent from a mobile handset to a cellular base station.
Usage-Based Insurance (UBI)	UBI bases insurance rates on pre-defined variables, including distance, driving behavior, time, and place.

V

Vehicle-to-Infrastructure (V2I)	The communication of smart cars and commercial vehicles with surrounding sensors.
Vehicle-to-Vehicle (V2V)	V2V systems allow vehicles to communicate with each other. Networks of vehicles can help avoid congestion, find better routes, and aid law enforcement.
Vehicle-to-Vehicle Communication (V2V Communication)	The wireless transmission of data between motor vehicles.
Vibration Sensing	A MEMS concept referring to the detection of periodic acceleration and deceleration. Typical applications include structural health monitoring, acoustic event triggering, and seismic equipment.
Virtual Private Network (VPN)	A secure system for users to send and receive data across shared or public networks. This is accomplished through encryption or protocols that act as if the user's devices were connected directly to the private network.
Virtual Sensor	Virtual sensors gather information that would not be measurable by a single device. This way, they can attain information that can't be measured directly.
Visited Location Register (VLR)	The database containing information about a subscriber's roaming within a mobile switching center's location area.

W

Wearable Technology

Technologies or computers integrated into articles of clothing or accessories that can be worn. Often, the wearable tech is used to quantify a physical process (such as heartbeat monitoring) or to augment human capabilities. Because of the impracticality of wires to transmit sensor data, wearables are almost universally wireless, using a variety of communication protocols, such as BLE.

Wide Area Network (WAN)

A telecommunications network or computer network that extends over a large geographical distance.

Wireless Application Protocol (WAP)

A protocol for wireless devices allowing the user to view and interact with data services. Often used to support internet access and Web browsing on mobile phones.

Wireless Fidelity (Wi-Fi)

A common form of local area network that operates on the 2.4 GHz band. Its popularity has led to a wide variety of devices to become Wi-Fi enabled, including smartphones, cameras, vehicles, and household appliances.

Z

ZigBee

Small-range wireless networking protocol that operates primarily on the 2.4 GHz frequency spectrum. ZigBee devices connect in a mesh topology, forwarding messages from controlling nodes to slaves, which repeat commands to other connected nodes. Due to its low power consumption and low data rate, ZigBee has been used in applications such as traffic management, wireless light switches, and industrial device monitoring.

Z-Wave

Wireless communication technology used in security systems, businesses, and home automation.



ABOUT THE AUTHOR

Syed Zaeem Hosain | Chief Technical Officer

Mr. Hosain is responsible for the architecture and future direction of Aeris' networks and technology strategy. He joined Aeris in 1996 as Vice President, Engineering and is a member of the founding executive team of Aeris. Mr. Hosain has more than 38 years of experience in the semiconductor, computer, and telecommunications industries, including product development, architecture design, and technical management.

Prior to joining Aeris, he held senior engineering and management positions at Analog Devices, Cypress Semiconductor, CAD National, and ESS Technology. Mr. Hosain is Chairman of the International Forum on ANSI-41 Standards Technology (IFAST) and Chairman of the IoT M2M Council (IMC). He holds a Bachelor of Science degree in Computer Science and Engineering from the Massachusetts Institute of Technology, Cambridge, MA.

ABOUT AERIS

Aeris is a global technology partner with a proven history of helping companies unlock the value of IoT. For more than a decade, we've powered critical projects for some of the most demanding customers of IoT services. Aeris strives to fundamentally improve businesses by dramatically reducing costs, accelerating time-to-market, and enabling new revenue streams. Built from the ground up for IoT and road tested at scale, Aeris IoT Services are based on the broadest technology stack in the industry, spanning connectivity up to vertical solutions. As veterans of the industry, we know that implementing an IoT solution can be complex, and we pride ourselves on making it simpler.

