White Paper

# What the Internet of Things (IoT) Needs to Become a Reality

Abstract

We're entering a new era of computing technology that many are calling the Internet of Things (IoT). Machine to machine, machine to infrastructure, machine to environment, the Internet of Everything, the Internet of Intelligent Things, intelligent systems—call it what you want, but it's happening, and its potential is huge.

We see the IoT as billions of smart, connected "things" (a sort of "universal global neural network" in the cloud) that will encompass every aspect of our lives, and its foundation is the intelligence that embedded processing provides. The IoT is comprised of smart machines interacting and communicating with other machines, objects, environments and infrastructures. As a result, huge volumes of data are being generated, and that data is being processed into useful actions that can "command and control" things to make our lives much easier and safer—and to reduce our impact on the environment.

The creativity of this new era is boundless, with amazing potential to improve our lives. What does the IoT need to become a reality? In this white paper, Freescale and ARM partner to answer that question.

Table of Contents

freescale™

ARM®
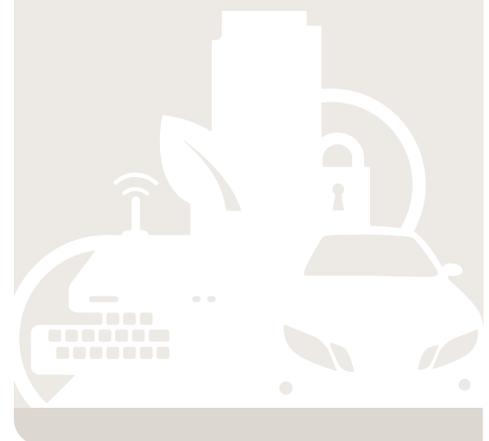
Introduction

Depending on who you talk to, the Internet of Things (IoT) is defined in different ways, and it encompasses many aspects of life—from connected homes and cities to connected cars and roads (yes, roads) to devices that track an individual's behavior and use the data collected for "push" services. Some mention one trillion Internet-connected devices by 2025 and define mobile phones as the "eyes and ears" of the applications connecting all of those connected "things." Depending on the context, others give examples that are less phone-centric, speak of a class of devices that do not exist today or point to Google's augmented-reality smart glasses as an indication of things to come.
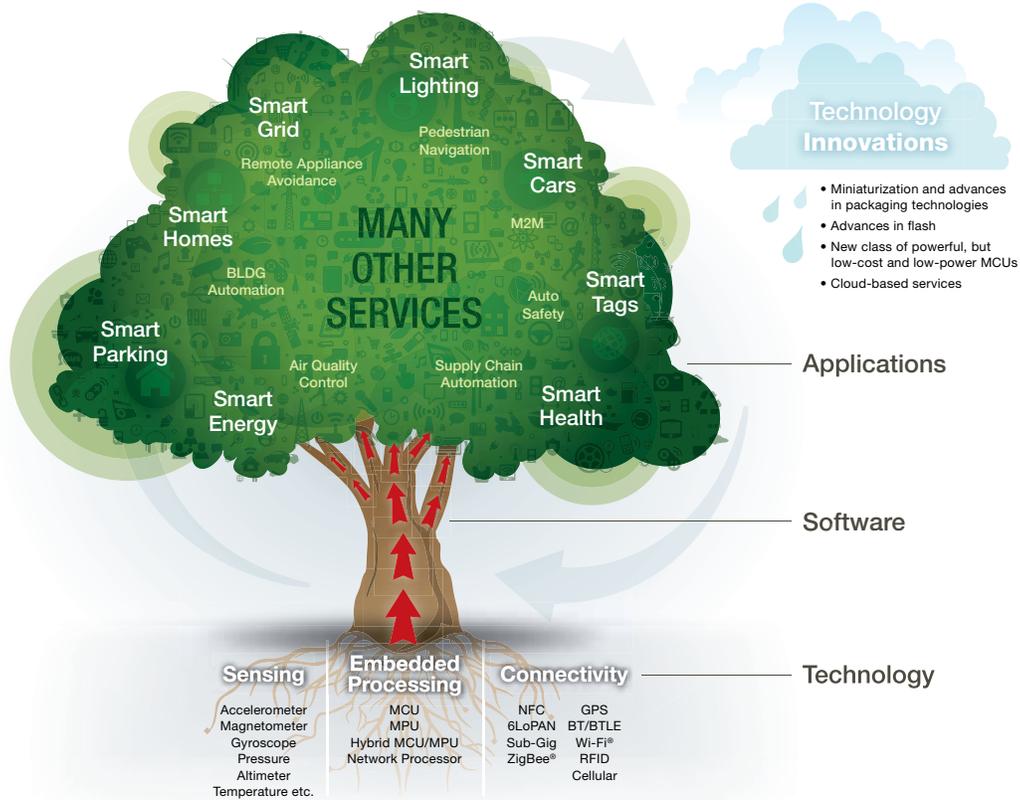
Everyone, however, thinks of the IoT as billions of connections (a sort of "universal global neural network" in the cloud) that will encompass every aspect of our lives. All of this public discussion suggests the IoT is finally becoming a hot topic within the mainstream media. Many recent articles point to the IoT as the interaction and exchange of data (lots of it) between machines and objects, and now there are product definitions reflecting the same concept. Hence, from a technology perspective, the IoT is being defined as smart machines interacting and communicating with other machines, objects, environments and infrastructures, resulting in volumes of data generated and processing of that data into useful actions that can "command and control" things and make life much easier for human beings … similar to the world envisioned in the 1970s cartoon *The Jetsons*, only better.

Estimates of the future market size of the IoT cover a broad range, but most pundits agree it will dwarf any other market. In mature markets today, the ultimate, pervasive consumer device is a mobile phone. Consider your own household, and count the number of mobile phones you currently have. Then count the number of windows, doors, electrical outlets, lights, appliances and heating and AC units you have. You'll quickly see why the IoT market will surpass the mobile phone market, at least in the western world.

A quick Internet search highlighted the following example use cases/applications under consideration:

- Machine-to-machine communication

- Machine-to-infrastructure communication

- Telehealth: remote or real-time pervasive monitoring of patients, diagnosis and drug delivery

- Continuous monitoring of, and firmware upgrades for, vehicles

- Asset tracking of goods on the move

- Automatic traffic management

- Remote security and control

- Environmental monitoring and control

- Home and industrial building automation

- "Smart" applications, including cities, water, agriculture, buildings, grid, meters, broadband, cars, appliances, tags, animal farming and the environment, to name a few

The IoT: Different Services, Technologies, Meanings for Everyone

Technology Innovations
- Miniaturization and advances in packaging technologies
- Advances in flash
- New class of powerful, but low-cost and low-power MCUs
- Cloud-based services

Applications

Software

Technology

Sensing
Accelerometer
Magnetometer
Gyroscope
Pressure
Altimeter
Temperature etc.

Embedded Processing
MCU
MPU
Hybrid MCU/MPU
Network Processor

Connectivity
NFC   GPS
6LoPAN   BT/BTLE
Sub-Gig   Wi-Fi®
ZigBee®   RFID
               Cellular

Making Things Smart

Do an IoT-related web search, and you'll quickly notice the overuse of the term "smart." So, what does it really mean when something is smart, and what makes an object smart? For example, how would a refrigerator or a toaster oven that hasn't been considered smart become a smart appliance?

Today, we are seeing the electrification of the world around us. Almost any manufactured good now includes an embedded processor (typically a microcontroller, or MCU), along with user interfaces, that can add programmability and deterministic "command and control" functionality. The electrification of the world and the pervasiveness of embedded processing are the keys to making objects "smart."

Your old toaster that mechanically controlled the color of your toast now has an MCU in it, and the MCU controls the color of your toast. The toaster completes its task more consistently and reliably, and because it is now a smart toaster, it has the ability to communicate with you electronically using its touchpad or switches.

After a device becomes smart through the integration of embedded processing, the next logical step is remote communication with the smart device to help make life easier. For example, if I am running late at the office, can I turn on my house lights for security reasons using my laptop or mobile phone?

Communication capability and remote manual control lead to the next step … how do I automate things and, based on my settings and with sophisticated cloud-based processing, make things happen without my intervention? That's the ultimate goal of some IoT applications. And, for those applications to connect with and leverage the Internet to achieve this goal, they must first become "smart" (incorporate an MCU/embedded processor with an associated unique ID) then connected and, finally, controlled. Those capabilities can then enable a new class of services that makes life easier for their users.

For the network, sophisticated cloud-based processing requires a new generation of communications processors that can keep track of all of those connected devices, communicate with them and translate their functionality into useful services … all with non-linear improvement to their performance and efficiency. The challenge will be to build secure networks that keep up with demand, while simultaneously reducing energy consumption and cost of equipment. This will require all kinds of innovations, well beyond the improvements Moore's law can deliver.

Application Categories

Let's look at some categories for IoT-related applications. While there are literally hundreds of applications being considered and identified by different industries, they can be categorized in a simple, logical way.

Category One

Category one encompasses the idea of millions of heterogeneous "aware" and interconnected devices with unique IDs interacting with other machines/objects, infrastructure, and the physical environment. In this category, the IoT largely plays a remote track, command, control and route (TCC&R) role. As with all aspects of the IoT, safety and security are paramount. These applications are not about data mining of people's behaviors (along the lines of "big brother watching") but rather they extend the automation and machine-to-machine (M2M), machine-to-infrastructure (M2I) and machine-to-nature (M2N) communications that can help simplify people's lives.

Category Two

The second category is all about leveraging the data that gets collected by the end nodes (smart devices with sensing and connectivity capability) and data mining for trends and behaviors that can generate useful marketing information to create additional commerce.

Credit card companies and membership shopping clubs already track and use people's behavior, to an extent, to come up with offers that may promote incremental sales. Now, the question is how far will this data mining go? Use cases could include a store tracking which aisles you visited, where you spent the most time within those aisles and even what type of items you lifted and browsed. This scenario is easily possible using a mobile phone's GPS capability, RFID and smart tags in stores and wireless tags. The result could be as simple as providing email offers or "push" services at the point of sale. Or, it could go further, with your car insurance company tracking your driving habits and places traveled to assign risk factors that help determine your monthly premium, for example. You can see how this category can become a slippery slope and how the IoT can enable data collection in every aspect of one's everyday life and assign a "category" to a person … with pleasant or unpleasant consequences.

When others become aware of the context associated with an entity, a person or a group (hence, knowing identity, location, activity and time), to what extent can that data be used, and to what extent should the entity, person or group have a say in how that data gets used? This second category, especially, spurs discussions about privacy, security, governance and the social responsibility that comes along with such a "self-aware," connected world.

This paper is focused on category one—specifically, the technologies and devices required to enable the IoT for TCC&R purposes.

IoT Use Cases

When devices can sense and communicate via the Internet, they can go beyond local embedded processing to access and take advantage of remote super-computing nodes. This allows a device to run more sophisticated analyses, make complex decisions and respond to local needs quickly, often with no human intervention required.

Let's take a look at the most common use cases for the IoT.

Pervasive Remote Tracking/Monitoring and (if needed) Command, Control and Routing (TCC&R)

This refers to remote tracking/monitoring and, if needed, command, control and routing functions for tasks and processes today usually done manually, or, if done remotely, that require additional infrastructure. For example, in most homes today, it's a manual process to turn on and off certain lights, set temperature zones and turn on and off a washing machine. In the future, doors, windows, electrical outlets, appliances and many other types of standalone equipment will become "smart" with a unique ID. Those smart devices can then be connected via wired or wireless communication, allowing a user to monitor his or her house remotely, change settings on a refrigerator or washing machine and control household tasks through a laptop or mobile phone. In fact, there are some services offered today by security or Internet service providers to do exactly that, but on a much smaller scale and with fewer capabilities than we expect to see in the future.

Remote Patient Monitoring
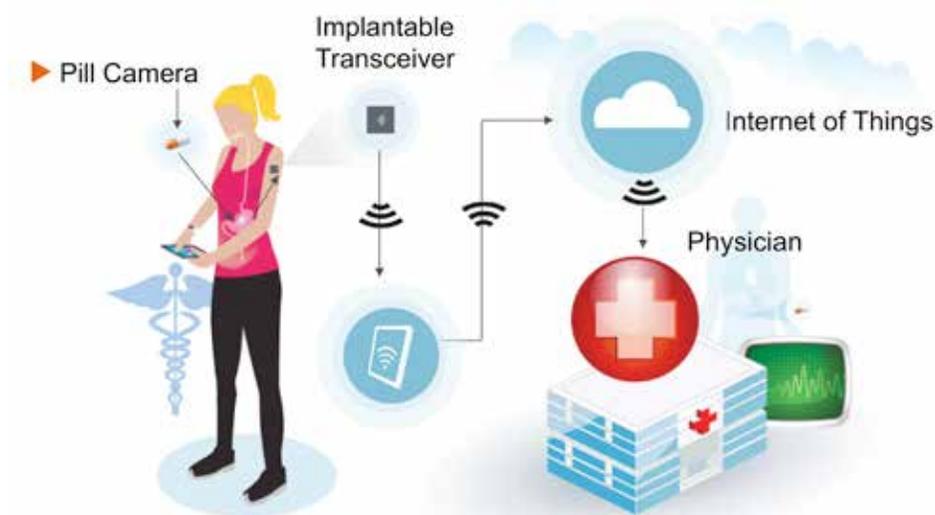
Asset Tracking

An extension of these kinds of services is asset tracking, which today is done via barcode and a variety of manual steps, but in the future will leverage smart tags, near-field communication (NFC) and RFID to globally track all kinds of objects, interactively. The word geo-tagged is now being used by some companies to refer to this class of applications. In a future scenario, a user would be able to use Google Earth to track anything with an RFID tag. Alternatively, your refrigerator could keep track of your smart-tagged groceries and tell your cell phone app you are low on a certain item. If your bag of frozen vegetables can have a smart tag, other objects such as valuable cars, jewelry and handbags could too, and they could be tracked via the Internet and also take advantage of a variety of available web-based applications.

Some telehealth-related services also belong in this category.

Process Control and Optimization

This is when various classes of sensors (with or without actuation capabilities) are used for monitoring and to provide data so a process can be controlled remotely. This could be as simple as the use of cameras (the sensing nodes in this example) to position boxes of various sizes on a conveyer belt so a label machine can properly apply labels to them. This task can be done in real time by sending the data to a remote computer, analyzing it and bringing a command back to the line so various control actions can be taken to improve the process … without any human intervention.

Resource Allocation and Optimization

The smart energy market provides an ideal example of this use case. The term "smart energy" has been used in many ways, but it basically refers to accessing information about energy consumption and reacting to the information to optimize the allocation of resources (energy use). In the case of a household, for example, once the residents know they've been using their washing machine during peak hours when the grid is most constrained and the cost of electricity is at premium, they could adjust their behavior and wash their laundry during non-peak hours, saving money and helping the utility company cope with the peak demand.

Context-aware Automation and Decision Optimization

This category is the most fascinating, as it refers to monitoring unknown factors (environmental, interaction between machines and infrastructures, etc.) and having machines make decisions that are as "human-like" as possible … only better!

Here's an example that can help illustrate this: "In a traffic collision avoidance system (TCAS), when two airplanes approach each other on a collision path, the 'machines' in the two airplanes take over. The system first sends an audible warning to the pilots about the danger ahead, while at the same time communicating between the two planes and deciding how each plane should move to avoid a collision. The assumption is that if the two pilots are warned and are in control to make quick decisions, they can both decide to make turns that would still cause a crash."

There are a whole host of new technologies available today and in development that could allow vehicles to communicate with each other as well as with a central control unit. These smart vehicles also could sense the road, traffic signs and lane markers and, using GPS and a communication link, avoid incoming traffic, avoid accidents around a curve or, in conjunction with the central control unit, avoid going over a distressed bridge on the verge of collapse.

Remote patient monitoring is another example relevant to this use case. For instance, imagine an implantable sensing node that tracks biometrics and sends a signal regarding an abnormal readout for an elderly patient. If the patient doesn't respond by taking a medication, the node could place an emergency call to a contact from a list, and, if there's no answer, call a second contact, and finally, if no answer, contact a monitoring clinic or quickly provide other emergency assistance. Another example is continuous monitoring of chronic diseases to help doctors determine best treatments, with minimal human intervention.

Requirements common to all of the use cases above include:

1) Sensing and data collection capability (sensing nodes)

2) Layers of local embedded processing capability (local embedded processing nodes)

3) Wired and/or wireless communication capability (connectivity nodes)

4) Software to automate tasks and enable new classes of services

5) Remote network/cloud-based embedded processing capability (remote embedded processing nodes)

6) Full security across the signal path

Functional View of Internet of Things Technologies
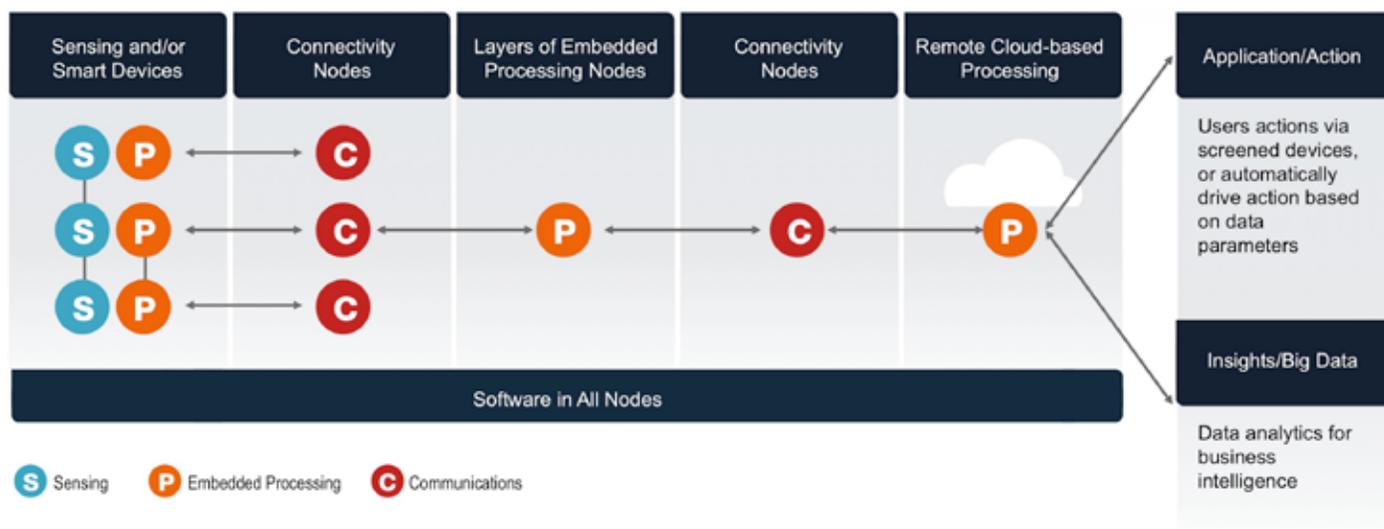
In the factory automation example (applying labels to boxes), a camera detects information using a charge-coupled device (CCD) sensor (sensing node), the collected data is then communicated to an embedded processor/controller (embedded processing node) using wired or wireless communication technology (connectivity node), a decision is made by the remote server (remote embedded processing node) and communicated (connectivity node), which causes a mechanical action to take place that corrects the situation.

A context-aware automation and decision optimization example could be a smart car using its active safety radar system (sensing node) in conjunction with image processing cameras (sensing nodes) that communicates with an embedded processor (embedded processing node) in the center stack of the car to make an appropriate decision regarding danger ahead. Or, the vehicle could leverage its built-in GPS and wide-area-network (WAN) wireless communication capability (connectivity node) to pass along information to a central processing server on the network/in the cloud (remote embedded processing node) that could then make the car aware of the information it had just received from the sensors on a bridge (sensing node) that was being pounded by flood waters and losing its structural integrity, guiding the car to a different route to avoid danger.

Building Blocks of the IoT

Sensing Nodes

The types of sensing nodes needed for the IoT vary widely, depending on the applications involved. Sensing nodes could include a camera system for image monitoring; water or gas flow meters for smart energy; radar vision when active safety is needed; RFID readers sensing the presence of an object or person; doors and locks with open/close circuits that indicate a building intrusion; or a simple thermometer measuring temperature. The bottom line is that there could be many different types of sensing nodes, depending on the applications. Who could forget the heat-seeking mechanical bugs that kept track of the population of a building in the movie *Minority Report*? Those mechanical bugs represent potential sensing nodes of the future.

These nodes will all carry a unique ID and can be controlled separately via a remote command and control topology. Use cases exist today in which a smartphone with RFID and/or NFC and GPS functionality can approach individual RFID/NFC-enabled "things" in a building, communicate with them and register their physical locations on the network. Hence, RFID and NFC will have a place in remote registration, and, ultimately, command and control of the IoT.

Layers of Local Embedded Processing Nodes

Embedded processing is at the heart of the IoT. Local processing capability is most often provided by MCUs, hybrid microcontrollers/microprocessors (MCUs/MPUs) or integrated MCU devices, which can provide the "real-time" embedded processing that is a key requirement of most IoT applications. Use cases vary significantly, and fully addressing the real-time embedded processing function requires a scalable strategy (using a scalable family of devices), as one size will not fit all.

In the home automation example, depending on the size or type of residence, requirements could vary from a simple network to a more complex structure with hierarchical, nested sub-networks controlled at different levels. For example, in a single-family home, all windows,

doors, electrical outlets and/or electrical equipment and thermostats could have simple embedded controllers that communicate with a master MCU/MPU hybrid device for command and control of the entire house. In turn, this master device can communicate via the Internet with a variety of "clients," from the security service provider and other service providers to portals that can give the homeowner access to remotely control all of these connected "things." In an apartment building, the same idea can be extended, with an even more complex, layered network hierarchy that includes apartment-level command and control, as well as floor-level and building-level command and control.

There are a few requirements that make an MCU ideal for use in the IoT.

- **Energy efficiency:** First and foremost, the MCU needs to be energy-efficient. In many cases, the sensing nodes are battery-operated satellite nodes, so a low-power spec is a basic requirement. For example, an MCU in a battery-operated thermostat that wakes up once every few minutes to check the temperature and adjust the AC based on its findings needs to consume as little power as possible to minimize battery replacement. Integrated circuit (IC) designers have many ways to reduce power consumption, including low-leakage process technologies, best-in-class low-power non-volatile memory/flash memory technologies, architectural innovations and various clocking schemes. For battery-operated nodes, all of those techniques are needed to achieve the lowest possible power consumption.

- **Embedded architecture with a rich software ecosystem:** The wide variety of potential IoT applications needs a software development environment that ties together the applications, the command, control and routing processing and the security of the node and system. While the importance of software in MCU solutions has increased during the past few years, for MCUs supporting the IoT, even more software, tools and enablement will be needed. A broad ecosystem with easily accessible support is key to enabling the development of embedded processing nodes and IoT applications.

- **Portfolio breadth that enables software scalability:** The ability to reuse software and leverage existing software investment is a key success factor for companies developing IoT applications. Software reuse enables the rapid rollout of multi-layered architectures (in which the embedded processor is tasked with different layers and levels of tracking, command, control and routing functions).

- **Portfolio breadth that cost-effectively enables different levels of performance and a robust mix of I/O interfaces:** The diversity of things to be controlled in the IoT, along with the different use cases, the number of things in a micro-network, different levels of service required and different interfaces in a heterogeneous environment will lead to the need for different tiers of devices, with diverse I/Os required for the various applications. A "one size fits all" approach will not be cost- or performance-optimized enough to satisfy the needs of this market.

- **Cost-effectiveness:** As with any other market, mass adoption will not take place until a certain price point for the solutions is reached. Like all other systems, the overall cost is the sum of the parts of the system plus the cost of the services required for the system. The overall system cost must be affordable for the paradigm shift to take hold in everyday life, so product cost is a very relevant factor.

• **Quality and reliability:** Unlike your mobile phone, laptop or other electronic device that you may change every two years, product life cycles in the industrial market are at least 10-15 years. Even inside a home, certain devices, such as thermostats, aren't changed that often. When you add the automotive market to the mix, more stringent reliability requirements and harsh environmental conditions must be supported. Hence, quality, reliability and longevity requirements for these markets are keys to the success of the IoT paradigm shift.

Although shifting the bulk of heavy-duty data processing and analysis to remote super-computing nodes in the network cloud is available and allows the local nodes "live longer" (not become obsolete as fast), there is still a balance between how much local vs. remote processing will be needed. This is especially important for time-critical applications that prefer local processing.

• **Security:** For the local embedded processing node at the physical layer, there are a variety of cryptographic engines and security accelerators to support data encryption (e.g. DES, AES, etc.) and authentication (e.g. SHA, etc.). Additional layers of security software, as well as best practices related to boot-up routines, are among the variety of security approaches available.

Wired and Wireless Communication Capability

The role of the communication node is to transfer information gathered by the sensing nodes and processed by local embedded processing nodes to the destinations identified by the local embedded processing nodes. And, once the data is remotely processed and new commands are generated, the communication node brings back the new commands to the local embedded processing nodes to execute a task.

Sometimes this could be as simple as sensing a fridge door being left open based on energy use, and after analyzing the data, automatically closing the door via a mechanical mechanism or generating a warning for the homeowners' "home automation app." Or, it could be as sophisticated as communication to an autonomous vehicle to avoid an accident.

Today's Wireless Landscape
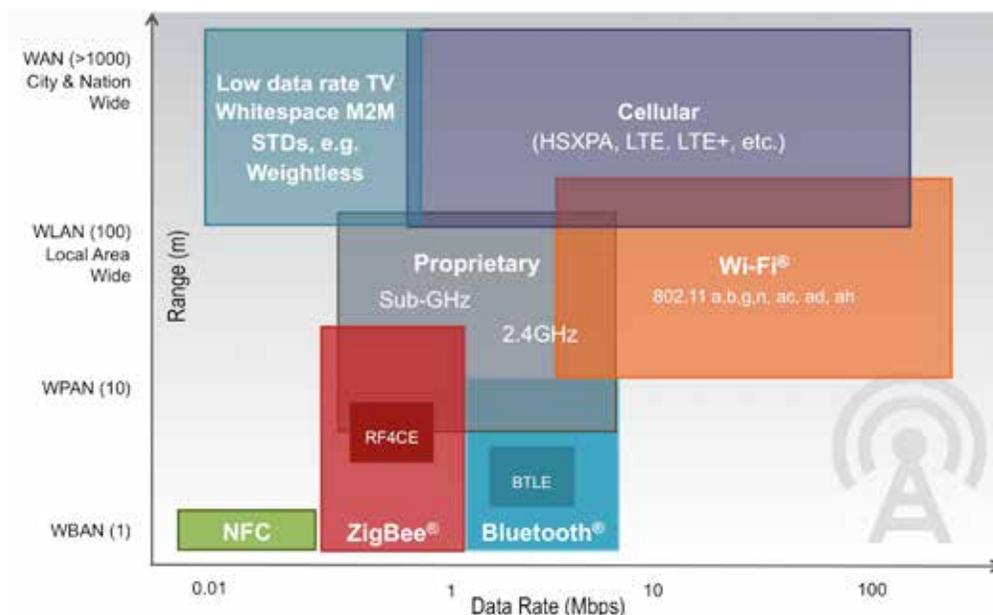
Use cases could vary drastically, but what is common to these command and control communication links is that they typically only need to carry few kilobytes of data for any given node, unless high-bandwidth image processing or video data is involved.

The IoT will encompass all aspects of one's everyday life, hence there is no limit to the distances for which command and control communication can/will be used. To get a better understanding of the dynamics of this segment, let's take a step back and look at the various communication topologies that exist today, from wireless body area network (WBAN) to wide area network (WAN), and all of the options in between. If you were to design wired and wireless technologies for the IoT from the ground up, you may or may not end up with the communications landscape as we know it today. However, many of the companies offering wireless and wired solutions are positioning their products as "the communication engine of choice" for the IoT market.

The IoT will also add the concept of wireless sensor and actuator networks (WSANs), which are networks that contain sensing and embedded processing nodes that can control their environment.

As with any emerging market, a transition period before system optimization takes place and technologies become better-suited for the end IoT-related applications is likely.

Based on typical product life cycles and the role of software, it would be safe to say that if a technology takes hold in an IoT segment now, that technology (or an optimized-to-purpose version of it) will be in place for at least the next five to eight years. There are some battle lines already drawn that may be solidifying. For example, it seems as though Bluetooth® Low Energy (BTLE) is being adopted by the health care industry for portable medical and lifestyle devices. On the other hand, the battle between ZigBee® and low-power Wi-Fi® technologies for industrial control and automation has just begun. Operators are urgently looking for new revenue streams, and machine-to-machine communication and location-based services seem to be good places to make a bet. Both can use existing infrastructure and are very much a part of the emerging IoT market.

Communication Technologies

	NFC	RFID	Blue-tooth®	Blue-tooth® LE	ANT	Proprietery (Sub-GHz & 2.4 GHz)	Wi-Fi®	ZigBee®	Z-wave	KNX	Wireless HART	6LoWPAN	WiMAX	2.5–3.5 G
Network	PAN	PAN	PAN	PAN	PAN	LAN	LAN	LAN	LAN	LAN	LAN	LAN	MAN	WAN
Topology	P2P	P2P	Star	Star	P2P, Star, Tree Mesh	Star, Mesh	Star	Mesh, Star, Tree	Mesh	Mesh, Star, Tree	Mesh, Star	Mesh, Star	Mesh	Mesh
Power	Very Low	Very Low	Low	Very Low	Very Low	Very Low to Low	Low-High	Very Low	Very Low	Very Low	Very Low	Very Low	High	High
Speed	400 Kbs	400 Kbs	700 kbs	1 Mbs	1 Mbs	250 kbs	11-100 Mbs	250 kbs	40 Kbs	1.2 Kbps	250 kbs	250 Kbs	11-100 Mbs	1.8-7.2 Mbs
Range	<10 cm	<3 m	<30 m	5-10 m	1-30 m	10-70 m	4-20 m	10-300 m	30 m	800 m	200 m	800 m (Sub-GHz)	50 km	Cellular network
Application	Pay, get access, share, initiate service, easy setup	Item tracking	Network for data exchange, headset	Health and fitness	Sports and fitness	Point to point connectivity	Internet, multimedia	Sensor networks, building and industrial automation	Residential lighting and automation	Building automation	Industrial sensing networks	Senor networks, building and industrial automation	Metro area broadband Internet connectivity	Cellular phones and telemetry
Cost Adder	Low	Low	Low	Low	Low	Medium	Medium	Medium	Low	Medium	Medium	Medium	High	High

Major volumes for the IoT market will likely not happen for another 10-12 years, and, at that time, the communications technologies may be completely different from those being considered today, or new revisions of existing standards may have emerged. Wi-Fi technologists already are working on 802.11ah (Wi-Fi on ISM bands below 1 GHz) to tailor it for infrastructure-independent ad-hoc, mesh networking and longer-range control of sensor networks. Alternatively, there could be brand-new technologies better suited for certain aspects of IoT communication that displace the existing standards for the IoT. For example, operators may decide their valuable spectrum is too precious to use for WAN-based command and control services and they instead need to use a different technology. Or, a disruptive wireless network technology like what Weightless (**weightless.org/**) is developing may take hold.

One thing about the connectivity needs of the future IoT market is clear—it is so diverse, large and cost-conscious that a range of different technologies will be needed (possibly including WAN, LAN, WPAN, WBAN, etc.), and one size will not fit all.

Requirements for communication functions are almost the same as for embedded processing nodes:

• Cost-effectiveness

• Low power

• Quality and reliability

• Security

'Box-level' View of IoT Building Blocks

If we convert the building blocks of the IoT from simple nodes to a box/product-level view, we end up with sensing/edge nodes that use PAN/BAN/LAN types of communications topologies, connected to gateways with different levels of hierarchy.

'Box-level' View of IoT Building Blocks
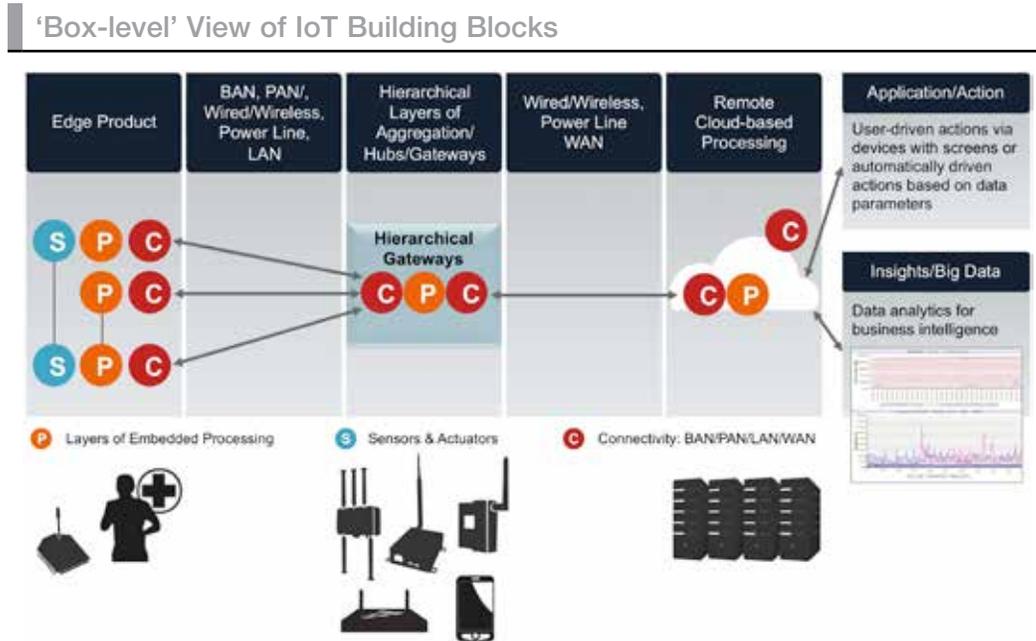
These gateways, in turn, communicate to the cloud via WAN communication technology. Once connected to the cloud through an access network, data will be routed through a server for application/action, as well as big data analysis.

Software Service Fabric for Metering Application
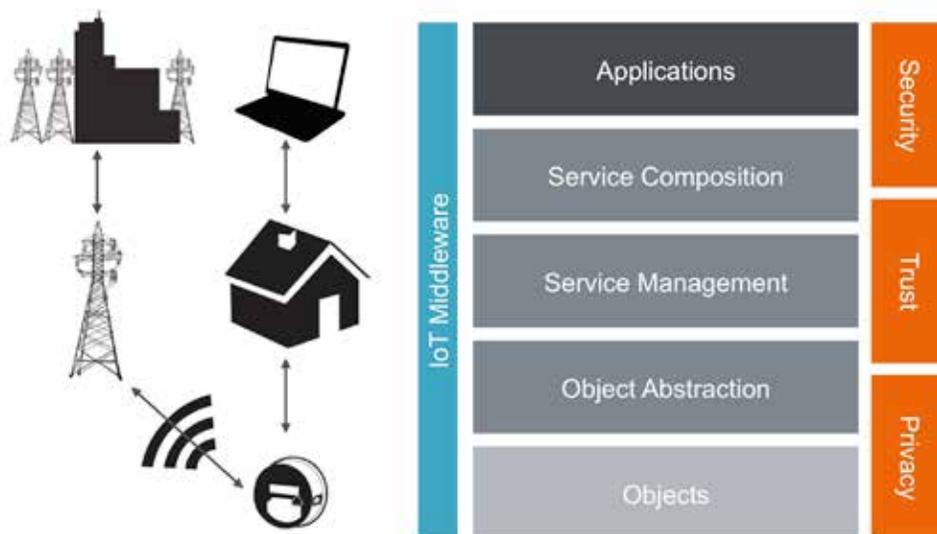
Software to Automate Tasks

Getting all segments of the IoT to communicate and work together is key to the success of the technology rollout, and that means deploying a lot of software (and middleware) that will enable various heterogeneous devices to talk with each other and the infrastructure around them.

For example, in a smart meter application, an analog front end (AFE) reads the meter and the MCU manages the meter to interpret and push the data through the communication pipe, which will be communicating with the house on one end and the curbside on the other end. While most developers have a clear view of the software architecture from a device, communication pipe and application profile perspective, the service-level fabric must also be considered for a given application.

In this configuration, the sensing node (here the AFE) is using an embedded processing (MCU) node to translate and transmit the data through the communication functions to the central embedded processing node in the house, as well as one on the curbside. A lot of middleware software is needed to enable this interaction to happen reliably, with the services delivered seamlessly.

Remote Embedded Processing Nodes (access to cloud computing)

Since there are not yet industry-wide IoT best practices agreed upon and deployed, many component providers are approaching the connection between devices and the cloud as a connection to their niche cloud, as opposed to the cloud. Some companies promote that all devices will be "dumb nodes," with all processing and decision-making done within "their cloud." Alternatively, some believe only minimal access to the cloud for basic Internet-related services will be required, with most of the "thinking" done locally. The architecture and building blocks of the IoT as described in this paper allow for a number of different approaches, which will likely be necessary due to the wide variety of use cases and configurations anticipated. That flexibility will be needed to optimize system-level performance.

So, why does software get such a big headline? Software enables the various services the IoT will provide. Services are the means by which the IoT will address certain needs. Those needs could exist today, or they may be things we don't yet realize we need, but someday we'll wonder why

we never had them before. Many people forget that until 20 years ago, most of us lived without mobile phones and didn't see a need for them, but now they are the most personal gadget owned by people in the western world. Along those lines, some IoT services will address needs easily identifiable today (e.g. asset tracking, smart energy, etc.), but others are yet to be defined.

Full Security Across the Entire Signal Path

Some people bundle this topic within the software portion of the IoT, but it deserves the attention of a separate category. Without a solid security mechanism for all of the IoT building blocks mentioned above, the IoT will not be as pervasive as it is anticipated to become.

When we say security, we really mean security of information—the information that gets passed around by various parts of the system and is context- and service-dependent. For example, knowing the location of a person could be considered a good thing if the person was lost. However, if that person felt his or her privacy was being compromised, knowing the location information could be considered a bad thing.

Here's what we mean by secure information:

- **Information needs to be available when needed:** This is the most basic level of security. If the information regarding an intruder in your house gets sent to the police station the next day, that information loses its value. The assurance that the services and their underlying infrastructure can process, store and deliver the data when and where it's needed is the first aspect of a secure system. In certain cases, redundant infrastructure needs is required to ensure this will happen.

- **Information needs to be confidential:** Hence, the owner of the information decides which authorized people, groups or organizations can access it. Safeguarding the information obtained by IoT services is critical, or those services will lose the users' trust. Mechanisms must be put in place to ensure confidentiality of the information exchanged. This is a tough balancing act, as there are a whole host of IoT-related services designed to leverage data mining and generate push services. The "opt out" mechanism for such services would be subject to the governance of the IoT.

- **The integrity of data needs to be assured:** Assurance that the information is accurate, authentic, timely and complete is key. Unless the data can be trusted and relied upon, it cannot be used for its intended purposes, and the entire service paradigm around that data will break down.

The security of the system is as good as the last threat it was able to prevent, and, as soon as it gets broken, one needs to implement new ways of making it secure again. If the recent hacking of credit card and personal information from reputable outlets on the Internet is any indication of the challenges facing IoT services, the Internet security infrastructure available today is inadequate to manage IoT services.

During the summer of 2010, malware targeted electronic process control systems for the first time instead of the traditional credit cards and personal information. The Stuxnet Trojan horse worm that attacked Siemens process control systems at nuclear plants demonstrated incredible levels of sophistication and showed the potential damage that could be done to undermine the security of the IoT.

Device-level Security: There are different types (MCU, hybrid MCU/MPU, integrated MCUs, etc.) and layers of embedded processing at various nodes of the IoT, and for any device to be considered smart so it can be connected to the Internet, it must incorporate an embedded processor. Embedded processors are going to be pervasive in the IoT, and they'd better be very secure.

MCUs are vulnerable during their boot-up process, when software is executed from programmable memory using the code stored in the read-only memory (ROM) or non-volatile memory (NVM)/flash memory. During this process, expert hackers can break the routine and hack the system in a variety of ways. Many new technologies are rolling out to address the security issues related to passive attacks (e.g. glitching) and invasive attacks (e.g. UV attacks), but more are likely necessary.

The intent of the IoT is to put smart devices on a sort of universal neural net, controlling them remotely. Hence, each of these identifiable objects (billions of them) can introduce a threat to the overall system. With such potential for disaster, are there best practices engineers can learn to enhance the security of MCUs in an IoT system?

By now it should be clear that networks of the future will connect more objects, machines and infrastructure to a global neural network of cloud-based services than they will connect people. A tsunami of data and services will affect the way we live, well beyond the changes experienced when the Internet first arrived and changed the way people network and communicate with each other. At the heart of the IoT are layers of embedded processing, from the most remote satellite sensing node to the core of the network. The diversity of services being planned for the IoT means no one company can develop full solutions and supporting IoT-based innovations. IoT-based innovations will require a broad, rich ecosystem of partner companies working together to bring IoT-based services to the market. An open (non-proprietary) platform (**ARM.com**) that allows all partners working together to use the same baseline technologies is key to making the IoT happen.

When Does the IoT Become a Reality?

The pervasiveness of embedded processing is already happening everywhere around us. At home, appliances as mundane as your basic toaster now come with an embedded MCU that not only sets the darkness of the piece of toast to your preference, but also adds functional safety to the device. Your refrigerator has started talking to you and keeping track of what you put in it. There are energy-aware HVAC systems that can now generate a report on the activity in your house and recommend ways to reduce your energy consumption. The electrification of vehicles has already started happening, and in just a few years from now, each car will contain >50 percent more electronics than it did just five years ago. The cars of the future will indeed be able to drive themselves. Similar changes are also happening in other aspects of our lives … in factories, transportation, school systems, stadiums and other public venues. Embedded processing is everywhere.

Connecting those smart devices (nodes) to the web has also started happening, although at a slower rate. The pieces of the technology puzzle are coming together to accommodate the Internet of Things sooner than most people expect. Just as the Internet phenomenon happened not so long ago and caught like a wildfire, the Internet of Things will touch every aspect of our lives in less than a decade. Are you ready for it?

Contributors

This white paper was developed in collaboration with Global Strategy and Business Development, Freescale and Emerging Technologies, ARM.

For more information about how Freescale and ARM are helping make the IoT a reality, visit freescale.com/IoT and arm.com.